arXiv:1608.02702v1 [cs.CV] 9 Aug 2016

Steerable Principal Components for Space-Frequency Localized Images

Boris Landa and Yoel Shkolnisky

August 10, 2016

Boris Landa
Department of Applied Mathematics, School of Mathematical Sciences
Tel-Aviv University
sboris20@gmail.com

Yoel Shkolnisky
Department of Applied Mathematics, School of Mathematical Sciences
Tel-Aviv University
yoelsh@post.tau.ac.il

Abstract

This paper describes a fast and accurate method for obtaining steerable principal components from a large dataset of images, assuming the images are well localized in space and frequency. The obtained steerable principal components are optimal for expanding the images in the dataset and all of their rotations. The method relies upon first expanding the images using a series of two-dimensional Prolate Spheroidal Wave Functions (PSWFs), where the expansion coefficients are evaluated using a specially designed numerical integration scheme. Then, the expansion coefficients are used to construct a rotationally-invariant covariance matrix which admits a block-diagonal structure, and the eigen-decomposition of its blocks provides us with the desired steerable principal components. The proposed method is shown to be faster then existing methods, while providing appropriate error bounds which guarantee its accuracy.

1 Introduction

Principal Component Analysis (PCA), also known as Karhunen-Loeve transform, is a ubiquitous method for dimensionality reduction which is often utilized for compression, de-noising, and feature-extraction from datasets. In its various forms, PCA serves as an important algorithmic building block in numerous applications in image processing and computer vision. Given a dataset, the basic idea behind classical PCA is to find the best linear approximation (in the least-squares sense) to the dataset using a set of orthonormal basis functions, thus allowing for a simple and effective processing method adaptive to the dataset at hand. Due to increasing improvements in acquisition and storage techniques, we often encounter the need to process very large datasets, which consist of many thousands of images of ever-growing resolutions. In addition, there exist datasets and acquisition techniques that introduce known deformation types during the process of obtaining the images, thus increasing the variability in the data.

In this work, we focus on the rotation deformation, and specifically, on the setting where each image was acquired through an unknown planar rotation. Therefore, it is only natural to include all planar rotations of all images when performing the PCA procedure. It is important to mention that when handling large datasets, the naive approach of introducing a large number of rotated versions of all images into the dataset, and then performing standard PCA, is computationally prohibitive, and moreover, is less accurate than considering the analytic continuum of all rotations. In the literature, one can find numerous works concerning the study of deformations (such as rotations, translations, dilations, etc.) and their connections to image processing and invariant-feature extraction through the theory of Lie groups and compact group theory [20, 4, 21, 22, 11]. In this context, we aim to incorporate the action of the group SO(2) into the framework of PCA of image datasets. We will refer to such a procedure as steerable-PCA. By the theory of Lie groups, and specifically the action of the group SO(2), it is known that the resulting principal components

(which best approximate all rotations of a set of images) are described by tensor products of radial functions and Fourier modes. Such functions are referred to as "steerable" [7], since they can be rotated (or "steered") by a simple multiplication with a complex-valued constant, and hence the term "Steerable" in Steerable-PCA. An early approach for computing the steerable-PCA was proposed in [25], which used the SVD to obtain the principal components of a single template and its set of deformations. In [9], the authors presented an efficient steerable-PCA algorithm based on Fourier decomposition of the images sampled on a polar grid. While this method considers the continuum of all planar rotations, it requires a non-obvious discretization of the images in the radial direction. Efficient methods for steerable-PCA were also introduced in [27] and [10]. These methods considered a finite set of equiangular rotations of each image on a polar grid, which allows for efficient circulant matrix decompositions to be carried out when computing the principal components. Recently, [37] and [36] utilized Fourier-Bessel basis functions to expand the images, followed by applying PCA in the domain of the expansion coefficients, thus accounting for all (infinitely many) rotations. We also mention [34], where the authors present an accurate algorithm to obtain the steerable principal components of templates whose analytical form is known in advance.

As digital images are typically specified by their samples on a Cartesian grid, considering their rotations implicitly assumes that they were sampled from some underlying continuous bivariate function. Rotation of a sampled image essentially interpolates this underlying continuous function from the available Cartesian samples. We remark that while previous works provide algorithms for steerable-PCA of discretized input images, they lack in describing rigorous connections between the results of the procedure and the continuous images prior to sampling, i.e. the underlying bivariate functions. In this work, we assume that the digital images were sampled from bivariate functions that are essentially bandlimited and are sufficiently concentrated in an area of interest in space. These assumptions guarantee that if an image was sampled with a sufficiently high sampling rate, it can be reconstructed from its samples with high precision [26]. Such assumptions are very common in various areas of engineering and physics, and as all acquisition devices are essentially bandlimited and restricted in space/time, are expected to hold for a wide range of image datasets. Since we are interested in processing images with arbitrary orientations, it is only natural to consider a circular support area, both in space and frequency, instead of the (classical) square. We note that this model was implicitly assumed to hold in [37] and [36] for Cryo Electron-Microscopy (Cryo-EM) images.

Under the model assumptions mentioned above, our goal is to develop a fast and accurate steerable-PCA procedure which considers the continuum of all planar rotations of all images in our dataset. Particularly appealing basis functions for expanding bandlimited functions which are also concentrated in space, are the Prolate Spheroidal Wave Functions (PSWFs) [32, 16, 17, 31, 24], defined as the strictly bandlimited set of orthogonal functions, which maximize the ratio between their $\mathcal{L}^2$ norm inside some finite region of interest and their $\mathcal{L}^2$ norm over the entire Euclidean

space. Recently, [13] described an approximation scheme for functions localized to disks in space and frequency, using a series of two-dimensional PSWFs. We therefore incorporate the methods introduced in [13] into the framework of steerable-PCA, providing accurate, scalable and efficient algorithms. Our approach is in the spirit of [36], resulting in a similar block-diagonal covariance matrix. However, replacing the Fourier-Bessel with PSWFs turns out to be advantageous in terms of accuracy, available error bounds, speed, and statistical properties.

The contributions of this paper are the following. By utilizing theoretical and computational tools related to PSWFs, we are able to provide accuracy guarantees for our steerable-PCA algorithm, under the assumptions of space-frequency localization. This accuracy is in part related to a rigorous truncation rule we provide for the PSWFs series expansion, in contrast to the series truncation rules used in [37] and [36]. In addition, using a quadrature integration scheme optimized for integrating bandlimited functions on a disk [30], we present an algorithm which is in theory (and in practice) faster than [36] by a factor between 2 and 4. Finally, we also show that under some conditions on the space-frequency concentration of the images at hand, the transformation to the PSWFs expansion coefficients, and the resulting principal components, are orthogonal to a very high precision.

The organization of this paper are as follows. In Section 2 we introduce the PSWFs and their usage in expanding space-frequency localized images specified by their equally-spaced Cartesian samples. In particular, we review the results of [13] on expanding a function into a series of PSWFs, evaluating the expansion coefficients, and bounding the overall approximation error. In Section 3 we present a fast algorithm for approximating the expansion coefficients from Section 2 up to an arbitrary precision. In Section 4 we formalize the procedure of steerable-PCA for the continuous setting (similarly to [34]), and combine it with our PSWFs-based approximation scheme. Section 5 then summarizes all relevant algorithms, and analyses in detail the computational complexities involved. In Section 6 we provide some numerical results on the spectrum of the transformation to the PSWFs expansion coefficients, which is of particular interest for noisy datasets, and in Section 7 we compare our algorithm to that of [36] in terms of running time and accuracy. Finally, in Section 8 we provide some concluding remarks and some possible future research directions.

2 Image approximation based on PSWFs expansion

Let $f : \mathbb{R}^2 \to \mathbb{R}$ be a square integrable function on $\mathbb{R}^2$. We define a function $f(x)$ as c-bandlimited if its two-dimensional Fourier transform, denoted by $F(\omega)$, vanishes outside a disk of radius c. Specifically, if we denote

$$\mathbf{D} \triangleq \{x \in \mathbb{R}^2,\ |x| \leq 1\},$$

then f is bandlimited with bandlimit c if

$$f(x) = \left(\frac{1}{2\pi}\right)^2 \int_{c\mathbf{D}} F(\omega) e^{\imath\omega x} d\omega, \quad x \in \mathbb{R}^2. \tag{1}$$

Among all c-bandlimited functions, the Prolate Spheroidal Wave Functions (PSWFs) on $\mathbf{D}$ (the unit disk), are the most energy concentrated in $\mathbf{D}$, i.e. the ratio between their energy inside $\mathbf{D}$ and their energy in all $\mathbb{R}^2$ is the largest. The 2D PSWFs were derived and analysed in [31], and they can be obtained as the solutions to the integral equation

$$\alpha\psi(x) = \int_{\mathbf{D}} \psi(\omega) e^{\imath c\omega x} d\omega, \quad x \in \mathbf{D}. \tag{2}$$

We denote the PSWFs with bandlimit c as $\psi^c_{N,n}(x)$, and their corresponding eigenvalues as $\alpha^c_{N,n}$, which together form the eigenfunctions and eigenvalues of (2), with N being a non-negative integer and n being a positive integer. In addition, it turns out that the functions $\psi^c_{N,n}(x)$ are orthogonal on both $\mathbf{D}$ and $\mathbb{R}^2$ using the standard $\mathcal{L}^2$ inner products on $\mathbf{D}$ and $\mathbb{R}^2$, respectively, and are respectively dense in both the class of $\mathcal{L}^2(\mathbf{D})$ functions and in the class of c-bandlimited functions on $\mathbb{R}^2$. Similarly to the Fourier-Bessel basis functions used in [37], the PSWFs have a separation of variables in polar coordinates, and can be written in complex form as

$$\psi^c_{N,n}(r,\theta) = \frac{1}{\sqrt{2\pi}} R^c_{N,n}(r) e^{\imath N\theta}, \quad N \in \mathbb{Z},\ n \in \mathbb{Z}^+, \tag{3}$$

where $R_{N,n}(r)$ (to be described shortly) satisfies $R_{N,n}(r) = R_{|N|,n}(r)$, and the eigenfunctions $\psi^c_{N,n}(x)$ are normalized to have an $\mathcal{L}^2(\mathbf{D})$ norm of 1. The indices N and n are often referred to as the angular index and the radial index respectively. Equation (3) also tells us that the PSWFs are steerable [7], that is, rotating $\psi^c_{N,n}(r,\theta)$ is equivalent to multiplying it by a complex constant. Therefore, every rotation of a linear combination of PSWFs remains a linear combination but with different coefficients, which depend on the rotation angle. This property is important for handling datasets which include rotations, and in particular, for the steerable-PCA procedure. Continuing with the evaluation of the PSWFs, the radial functions $R^c_{N,n}(r)$ in (3) are real-valued, and are obtained as the solutions to the integral equation

$$\beta R(r) = \int_0^1 R(\rho) J_N(cr\rho) \rho d\rho, \quad r \in [0,1], \tag{4}$$

where $J_N(x)$ is the Bessel function of the first kind of order N. For every index N, the functions $R^c_{N,n}(r)$ together with their associated eigenvalues $\beta^c_{N,n}$ form the eigenfunctions and eigenvalues of equation (4), where $N \in \mathbb{Z}$ and $n \in \mathbb{Z}^+$. The eigenvalues $\alpha^c_{N,n}$ and $\beta^c_{N,n}$ of (3) and (4) are related by

$$\alpha^c_{N,n} = 2\pi\imath^N \beta^c_{N,n}. \tag{5}$$

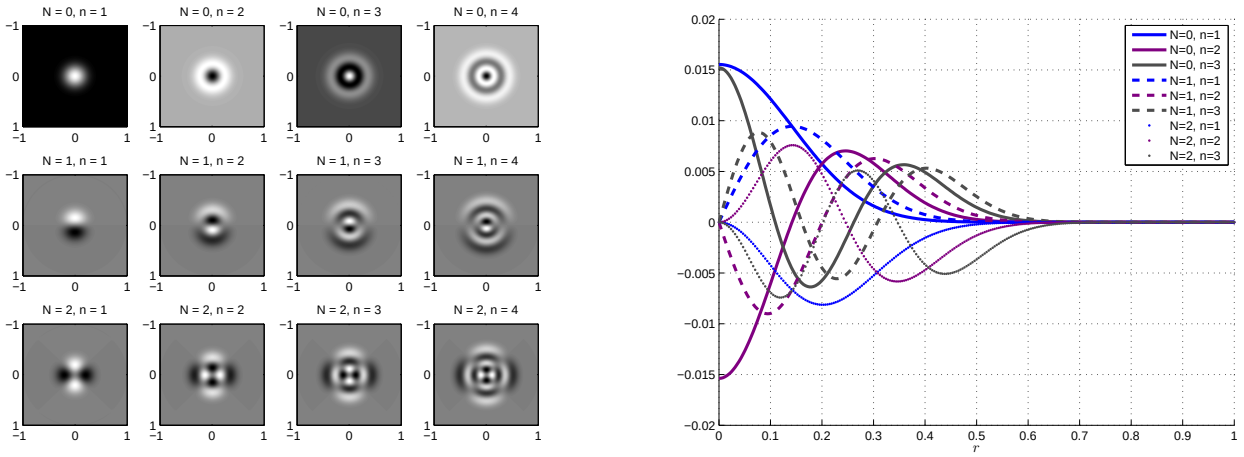

(a) Two-dimensional view

(b) Radial profile

Figure 1: Illustration of the first few PSWFs with bandlimit $c = 16\pi$, ordered according to their eigenvalues (for every index N, the eigenvalues are ordered from largest to smallest as a function of the index n).

A numerical evaluation procedure for the two-dimensional PSWFs can be found in [30], and an illustration of the PSWFs for the several first index pairs (N, n) can be seen in Figure 1. Since our images are assumed to be essentially bandlimited and sufficiently concentrated in a disk, the properties of the PSWFs mentioned above (and especially their optimal concentration property) make them suitable basis functions for expanding such images, as we consider next. Let us define a function $f(x)$ as (ν, μ)-concentrated if its $\mathcal{L}^2$ norm outside a disk of radius ν is upper bounded by μ, that is

$$\sqrt{\int_{x \notin \nu \mathbf{D}} |f(x)|^2 dx} \leq \mu. \tag{6}$$

It can be seen that the class of c-bandlimited functions is a subclass of (c, δ_c)-concentrated functions in the Fourier domain (for any $\delta_c > 0$). We consider the two-dimensional functions from which our images were sampled to be $(1, \varepsilon)$-concentrated (in space), and such that their Fourier transforms are (c, δ_c)-concentrated. This assumption always holds for some set of parameters c, δ_c, ε, and the general notion is that δ_c and ε are "small". Since our images are given in their sampled form, we define the unit square $\mathbf{Q} \triangleq [-1, 1] \times [-1, 1]$, and assume that we are given the samples $\left\{ f(x_k) : x_k = \frac{k}{L} \in \mathbf{Q} \right\}$, where k is a two-dimensional integer index. These samples correspond to a Cartesian grid of $(2L + 1) \times (2L + 1)$ equally spaced samples, with sampling rate of L in each dimension. We mention that in this setup, the Nyquist frequency corresponds to a bandlimit c of at most πL.

Let us suppose that we are given M (sampled) images. We denote the samples of the m'th image as $I_m^k \triangleq I_m(x_k)$, where $I_m(x)$ is the m'th continuous image, $0 \leq m \leq M - 1$, and $x_k \in \mathbf{Q}$.

If we assume that each continuous image $I_m(x)$ is in $\mathcal{L}^2(\mathbf{D})$, we can then expand it as

$$I_m(x) = \sum_{N=-\infty}^{\infty} \sum_{n=1}^{\infty} a_{N,n}^m \psi_{N,n}^c(x), \quad x \in \mathbf{D}. \tag{7}$$

As numerical algorithms cannot use infinite expansions, our first goal is to expand the images using only a finite sequence of PSWFs, that is

$$I_m(x) \sim \sum_{N,n \in \Omega} a_{N,n}^m \psi_{N,n}^c(x), \tag{8}$$

where Ω is a finite set of indices, c is our chosen bandlimit for the PSWFs, and the expansion coefficients $\left\{a_{N,n}^m\right\}_{m=0}^{M-1}$ are evaluated using the available samples $\left\{I_m^k\right\}_{m=0}^{M-1}$. To this end, we adopt the theorems and methods given in [13], where a PSWFs-based approximation scheme was introduced for expanding essentially bandlimited and space-concentrated functions on a disk. First, we address the issue of coefficients evaluation using a slightly modified version of the PSWFs, such that the appropriate expansion coefficients can be obtained using simple scalar products. Specifically, let

$$\hat{\psi}_{N,n}^c(x) \triangleq \left(\frac{c}{2\pi L}\right) \alpha_{N,n}^c \psi_{N,n}^c(x), \tag{9}$$

where $\alpha_{N,n}^c$ is the eigenvalue from (2) corresponding to the eigenfunction $\psi_{N,n}^c(x)$. We define the modified PSWFs expansion that approximates $I_m(x)$ as

$$\hat{I}_m(x) \triangleq \sum_{N,n \in \Omega} \hat{a}_{N,n}^m \hat{\psi}_{N,n}^c(x), \tag{10}$$

and the approximation coefficients as

$$\hat{a}_{N,n}^m \triangleq \left\langle I_m, \hat{\psi}_{N,n}^c \right\rangle = \sum_{\frac{k}{L} \in \mathbf{D}} I_m(\frac{k}{L}) \left[\hat{\psi}_{N,n}^c(\frac{k}{L})\right]^*, \tag{11}$$

where $(\cdot)^*$ denotes complex conjugation, and both I_m and $\hat{\psi}_{N,n}^c$ stand for the appropriate vectors of samples of $I_m(x)$ and $\hat{\psi}_{N,n}^c(x)$ inside the unit disk, respectively. We note that for real-valued images we have that $a_{N,n}^m = \left[a_{-N,n}^m\right]^*$, thus it is sufficient to compute the approximated coefficients $\hat{a}_{N,n}^m$ only for $N \geq 0$. Next, we describe a method for choosing Ω in (10). For any real positive T, define the index set

$$\Omega_T \triangleq \left\{ (N,n) : \sqrt{\frac{\left|\lambda_{N,n}^c\right|^2}{1 - \left|\lambda_{N,n}^c\right|^2}} > T \right\}, \tag{12}$$

where

$$\lambda_{N,n}^c \triangleq \frac{c}{2\pi} \alpha_{N,n}^c. \tag{13}$$

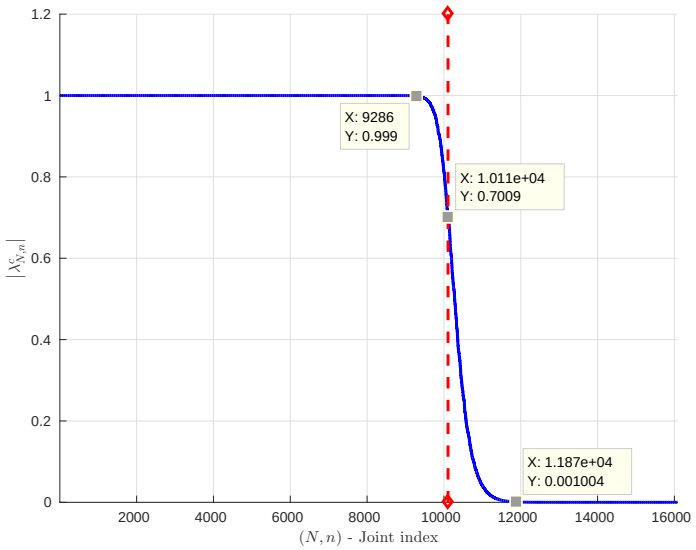

Figure 2: Illustration of the normalized eigenvalues $\left|\lambda^c_{N,n}\right|$ for $L = 64$, $c = \pi L$, sorted in a non-increasing order with a joint index k enumerating over (N, n). It is evident that for $T \ll 1$ the truncation will take place to the right of the vertical dashed line, where the normalized eigenvalues rapidly decay to zero. On the other hand, for $T \gg 1$ the truncation will take place to the left of the vertical line, where the normalized eigenvalues are extremely close to one. The dashed vertical line corresponds to $k = c^2/4$, which also approximately agrees with $\left|\lambda^c_{N,n}\right|^2 = 0.5$ (or equivalently with the truncation rule $T = 1$).

Figure 2 depicts the behaviour of $\left\{\left|\lambda^c_{N,n}\right|\right\}_{(N,n)}$ for $L = 64$ and $c = \pi L$.

Using our definitions in (10), (11), (12), together with our prior assumptions on the concentration properties of the images, it follows from the approximation theorems in [13], that the approximation error norm of the m'th continuous image using the truncated PSWFs expansion satisfies

$$\left\| I_m(x) - \sum_{N,n\in\Omega_T} \hat{a}^m_{N,n}\hat{\psi}^c_{N,n}(x) \right\|_{\mathcal{L}^2(\mathbf{D})} \leq \left(\varepsilon + \frac{\delta_c}{2\pi}\right) T + \eta\left(\frac{c}{2\pi L}\right)^2 \sqrt{\sum_{\frac{k}{L}\notin\mathbf{D}} \left|I_m(\frac{k}{L})\right|^2} + \frac{2}{\pi}\delta_c, \qquad (14)$$

with η being a constant no bigger than $2\pi^2 L/c$. Roughly speaking, the first term of the error bound in (14) results from the truncation of the series in (7), the second term results from the coefficients' evaluation error (since we only use a finite set of samples in the coefficients' evaluation), and the third term results from truncation of the images in the Fourier domain to the bandlimit c. We can somewhat simplify the bound in (14) by considering our images $I_m(x)$ to be restricted to the unit disk, thus eliminating the dependence of the error bound on their samples outside the unit disk.

Specifically, we show in Appendix A that the bound of (14) can be replaced with

$$\left\| I_m(x) - \sum_{N,n\in\Omega_T} \hat{a}^m_{N,n}\hat{\psi}^c_{N,n}(x) \right\|_{\mathcal{L}^2(\mathbf{D})} \leq \left(\varepsilon + \frac{\delta_c}{2\pi}\right)(T+4) \triangleq \mathcal{E}(\varepsilon,\delta_c,T), \tag{15}$$

which does not depend on the samples of the images outside the unit disk, but only on their global concentration properties together with the truncation parameter T.

Lastly, it is also mentioned in [13] that the cardinality of the index set Ω_T of (12) is expected to be

$$|\Omega_T| = \frac{c^2}{4} - \frac{2}{\pi^2} c\log(c)\log(T) + o(c\log(c)). \tag{16}$$

The first term on the right hand-side of (16), which depends quadratically on c, is in-fact (up to a constant of $1/\pi$) what is known as the Landau rate [14] for stable sampling and reconstruction, and it is the term which dominates asymptotically the required number of basis functions for the approximation (see also Figure 2). We remark that for $c = \pi L$ (Nyquist sampling) and $T = 1$, which inherently means that the approximation error is of the order of the space-frequency localization, we have that the number of chosen basis functions in Ω_T is about $\pi^2 L^2/4$, which is smaller than the number of samples (pixels) in the unit disk.

We note that in the approximation scheme described in this section, the evaluation of the expansion coefficients for every image I_m requires $O(L^4)$ operations. This is due to the fact that each image contains $O(L^2)$ samples, while there are about $O(L^2)$ different basis functions in the expansion. In what follows, we describe an asymptotically more efficient method for evaluating the expansion coefficients in $O(L^3)$ operations.

3 Fast PSWFs coefficients approximation

Let us look at the formula for the PSWFs expansion coefficients from (11)

$$\hat{a}^m_{N,n} = \left\langle I_m, \hat{\psi}^c_{N,n} \right\rangle = \sum_{\frac{k}{L}\in\mathbf{D}} I_m(\frac{k}{L}) \left[\hat{\psi}^c_{N,n}(\frac{k}{L})\right]^* = \left(\frac{c}{2\pi L}\right) \sum_{\frac{k}{L}\in\mathbf{D}} I_m(\frac{k}{L}) \left[\alpha^c_{N,n}\psi^c_{N,n}(\frac{k}{L})\right]^*, \tag{17}$$

and rewrite it using (2) as

$$\begin{aligned} \hat{a}^m_{N,n} &= \left(\frac{c}{2\pi L}\right) \sum_{\frac{k}{L}\in\mathbf{D}} I_m(\frac{k}{L}) \left[\int_{\mathbf{D}} \psi^c_{N,n}(\omega) e^{\imath c\omega\frac{k}{L}} d\omega\right]^* \\ &= \left(\frac{c}{2\pi L}\right) \int_{\mathbf{D}} \left[\psi^c_{N,n}(\omega)\right]^* \sum_{\frac{k}{L}\in\mathbf{D}} I_m(\frac{k}{L}) e^{-\imath c\omega\frac{k}{L}} d\omega. \end{aligned} \tag{18}$$

In addition, let us define the function

$$\phi_m^c(u) \triangleq \sum_{\frac{k}{L}\in\mathbf{D}} I_m(\frac{k}{L}) e^{-\imath c u \frac{k}{L}}, \quad u \in \mathbb{R}^2. \tag{19}$$

Since both $\psi_{N,n}^c$ and ϕ_m^c are c-bandlimited functions, the product $\psi_{N,n}^c \cdot \phi_m^c$ defines a function which is $2c$-bandlimited. Therefore, we can compute (18) using a quadrature formula for $2c$-bandlimited functions on a disk. Such integration scheme was proposed in [30] using specialized quadrature nodes and weights. Thus, using an appropriate quadrature formula we can write the integral in (18) as

$$\int_{\mathbf{D}} \left[\psi_{N,n}^c(\omega)\right]^* \phi_m^c(\omega) d\omega = \sum_{\ell=1}^{\mathcal{N}_r} \sum_{j=1}^{\mathcal{N}_\theta^\ell} \mathcal{W}_{\ell,j}^{2c} \left[\psi_{N,n}^c(\omega_{\ell,j}^{2c})\right]^* \phi_m^c(\omega_{\ell,j}^{2c}) + \epsilon_q, \tag{20}$$

where $\omega_{\ell,j}^{2c}$ and $\mathcal{W}_{\ell,j}^{2c}$ are the quadrature nodes and weights respectively, $\mathcal{N}_r$ is the number of different radii of the quadrature nodes, $\mathcal{N}_\theta^\ell$ is the number of quadrature nodes per radius (which may vary for each radius), and ϵ_q is the remaining numerical integration error. We use for $\omega_{\ell,j}^{2c}$ and $\mathcal{W}_{\ell,j}^{2c}$ the quadrature nodes and weights of [30] corresponding to bandlimit $2c$. These quadrature nodes are equally-spaced in the angular direction and non-equally spaced in the radial direction. Thus, by using polar coordinates, we can describe the quadrature nodes as

$$\omega_{\ell,j}^{2c} = \left(r_\omega^\ell, \theta_\omega^{\ell,j}\right), \tag{21}$$

where the angles $\theta_\omega^{\ell,j}$ are given explicitly by the formula

$$\theta_\omega^{\ell,j} = \frac{2\pi j}{\mathcal{N}_\theta^\ell}. \tag{22}$$

Moreover, the quadrature weights from (20) depend only on the radius of their associated nodes, and explicitly,

$$\mathcal{W}_{\ell,j}^{2c} = \frac{2\pi}{\mathcal{N}_\theta^\ell} r_\omega^\ell \tilde{\mathcal{W}}_\ell^{2c}, \tag{23}$$

where the derivation of $\tilde{\mathcal{W}}_\ell^{2c}$ can also be found in [30]. By employing all of the above, together with the analytical expression of the PSWFs (3), we finally arrive at

$$\begin{aligned}
\hat{a}_{N,n}^m &= \left(\frac{c}{2\pi L}\right) \int_{\mathbf{D}} \left[\psi_{N,n}^c(\omega)\right]^* \phi_m^c(\omega) d\omega \\
&= \frac{c}{\sqrt{2\pi} L} \sum_{\ell=1}^{\mathcal{N}_r} \tilde{\mathcal{W}}_\ell^{2c} \frac{R_{N,n}^c(r_\omega^\ell) r_\omega^\ell}{\mathcal{N}_\theta^\ell} \sum_{j=1}^{\mathcal{N}_\theta^\ell} \phi_m^c(\omega_{\ell,j}^{2c}) e^{-\imath N \theta} + \frac{c}{2\pi L} \epsilon_q,
\end{aligned} \tag{24}$$

which leads us to define our approximated expansion coefficients as

$$\tilde{a}_{N,n}^{m} \triangleq \frac{c}{\sqrt{2\pi}L} \sum_{\ell=1}^{\mathcal{N}_r} \tilde{\mathcal{W}}_{\ell}^{2c} \frac{R_{N,n}^{c}(r_{\omega}^{\ell}) r_{\omega}^{\ell}}{\mathcal{N}_{\theta}^{\ell}} \sum_{j=1}^{\mathcal{N}_{\theta}^{\ell}} \phi_{m}^{c}(\omega_{\ell,j}^{2c}) e^{-\imath N\theta}, \tag{25}$$

and the resulting numerical integration error as

$$\tilde{\epsilon}_q \triangleq \frac{c}{2\pi L} \epsilon_q. \tag{26}$$

We note that (25) is particularly appealing for efficient numerical evaluation, as we explain later in this section.

Next, we are interested in choosing appropriate values for $\mathcal{N}_r$ and $\mathcal{N}_{\theta}^{\ell}$ such that the numerical integration error $\tilde{\epsilon}_q$ is sufficiently small, i.e.

$$|\tilde{\epsilon}_q| = \frac{c}{2\pi L} |\epsilon_q| \leq |\epsilon_q| \leq \vartheta_q, \tag{27}$$

for $c \leq \pi L$ and with ϑ_q being a small constant (usually chosen of the order of machine precision). To this end, and according to [30], it is sufficient to satisfy

$$\sum_{|j|>\mathcal{N}_{\theta}^{\ell}} J_j(2cr_{\omega}^{\ell}\rho) \leq C_1 \vartheta_q, \tag{28}$$

for all $1 \leq \ell \leq \mathcal{N}_r$ and $\rho \in [0,1]$, as well as

$$\sum_{k=2\mathcal{N}_r+1}^{\infty} \frac{\left|\lambda_{0,k}^{2c}\right|}{2c} \left\|R_{0,k}^{2c}(r)\right\|_{\infty} \left\|R_{0,k}^{2c}(r)\sqrt{r}\right\|_{\infty} \left(1 + \sum_{j=1}^{\mathcal{N}_r} \left|\tilde{\mathcal{W}}_{j}^{2c}\right|\right) \leq C_2 \vartheta_q, \tag{29}$$

where $\lambda_{N,n}^{c}$ was defined in (13), $\|\cdot\|_{\infty}$ stands for the max-norm in $[0,1]$, and both C_1 and C_2 are constants which depend on the bandlimit c. We mention that $\left|\lambda_{0,k}^{2c}\right|$ are ordered in a non-increasing order with respect to the index k. In order to determine appropriate values for $\mathcal{N}_r$ and $\mathcal{N}_{\theta}^{\ell}$, one can proceed to solve the inequalities in conditions (28) and (29) numerically by directly evaluating these expressions. Plotted in Figure 3 is a typical array of quadrature nodes obtained via the scheme of [30].

In order to analyse the computational complexity of the procedure described in this section, and to offer the reader simpler analytic expressions and some insight concerning conditions (28) and (29), we provide some further analysis of these conditions and of the resulting number of quadrature nodes. In Appendix B, we prove that in order to satisfy condition (28) for every $1 \leq \ell \leq \mathcal{N}_r$, it is sufficient to choose the integers $\mathcal{N}_{\theta}^{\ell}$ as

$$\mathcal{N}_{\theta}^{\ell} = \lceil cr_{\omega}^{\ell} e + \log \vartheta_q^{-1} + \log \frac{2}{C_1} + 1 \rceil, \tag{30}$$

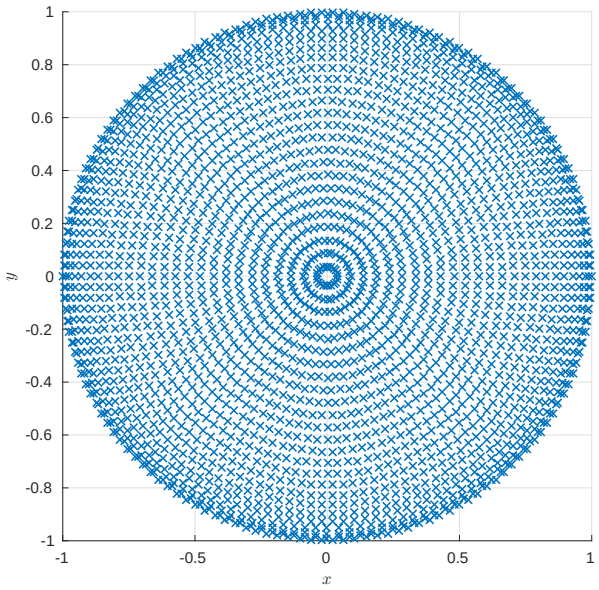

Figure 3: Quadrature nodes $\omega_{\ell,j}^{2c} = \left(r_\omega^\ell, \theta_\omega^{\ell,j}\right)$ for $L = 16$, $c = \pi L$, and accuracy of the order of machine precision. The nodes are arranged to be equally-spaced in the angular direction (the specific phase of the allocation is arbitrary) with different numbers of nodes for each radius (rings closer to the origin require less angular nodes). The number of radial nodes is about c/π, and the number of angular nodes per radius r_ω^ℓ is about $cr_\omega^\ell e$, where e is the base of the natural logarithm.

where r_ω^ℓ are the radial quadrature nodes from (21), e is the base of the natural logarithm, and $\lceil\cdot\rceil$ is the rounding up operation. It is evident that the term $cr_\omega^\ell e$ dominates condition (30), since all other factors (prescribed error ϑ_q and the constant C_1) affect it only logarithmically. Therefore, we expect the overall number of quadrature nodes to be

$$\sum_{\ell=1}^{\mathcal{N}_r} \mathcal{N}_\theta^\ell \sim ce \sum_{\ell=1}^{\mathcal{N}_r} r_\omega^\ell \sim \frac{ce\mathcal{N}_r}{2}, \tag{31}$$

assuming that the quadrature points r_ω^ℓ are approximately symmetric about 0.5. We remark that numerical experiments reveal that there is a small difference between condition (28) and (30), and specifically, choosing $\mathcal{N}_\theta^\ell$ directly via the numerical evaluation of condition (28) results in about 20% less angular nodes compared to choosing $\mathcal{N}_\theta^\ell$ according to (30) alone.

With regard to condition (29), although it may seem somewhat daunting, the sum on the left hand-side is dominated by the values of $\left|\lambda_{0,k}^{2c}\right|$ and their decay properties. We argue in Appendix C that the number of non-negligible (relative to machine precision) values of $\left|\lambda_{0,k}^{2c}\right|$ is only about $2c/\pi$, that is

$$\left|\left\{k : \ \left|\lambda_{0,k}^{2c}\right| > \epsilon\right\}\right| \sim \frac{2c}{\pi}, \tag{32}$$

for some small ϵ and a sufficiently large c. This observation stems from the fact that the values

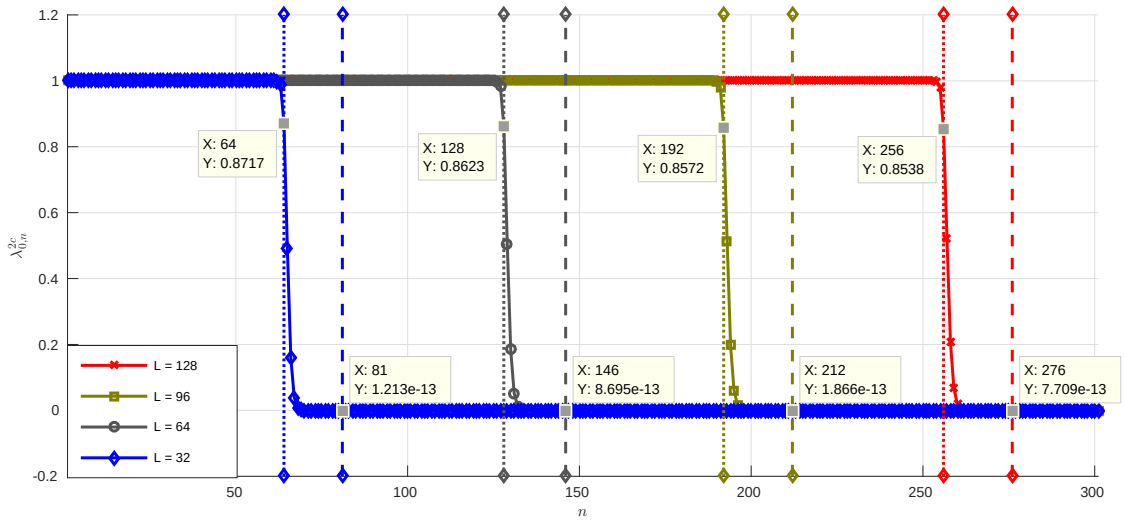

Figure 4: Behaviour of $\left|\lambda_{0,n}^{2c}\right|$ for $c = \pi L$ (Nyquist sampling) and several values of L. For every value of L, the dotted vertical line represents $n = 2c/\pi$, and the dashed vertical line represents the first index n for which $\left|\lambda_{0,n}^{2c}\right| \leq 10^{-12}$, which we denote by n_1^{2c}. First, we observe that $2c/\pi$ is the exact index for which $\left|\lambda_{0,n}^{2c}\right|$ begins its rapid decay, and second, it is evident that the difference between n_1^{2c} and the estimate $2c/\pi$ grows extremely slowly with L, making the asymptotic estimate $n_1^{2c} \sim 2c/\pi$ sufficiently accurate even for moderate values of L.

of $\left|\lambda_{0,k}^{2c}\right|$ become arbitrarily small (and decay at a super-exponential rate) once k reaches $2c/\pi + O(\log c)$. We provide numerical evidence for this claim in Figure 4.

Following the last observation, we therefore expect that the number of radial nodes $\mathcal{N}_r$ which is sufficient to satisfy condition (29) will be given for $c = \pi L$ (Nyquist sampling) by

$$\mathcal{N}_r \sim \frac{c}{\pi} = L. \tag{33}$$

Overall, using (31) and (33), the total number of quadrature nodes for $c = \pi L$ is

$$\sum_{\ell=1}^{\mathcal{N}_r} \mathcal{N}_\theta^\ell \sim \frac{e}{2\pi} c^2 = \frac{\pi e}{2} L^2, \tag{34}$$

making the number of quadrature nodes only slightly larger than the number of sampling points in the unit square, which is $(2L+1)^2$.

Let us now turn our attention to the computational complexity of computing $\tilde{a}_{N,n}^m$ of (25). The first step is to compute $\phi_m^c(\omega_{\ell,j}^{2c})$ from (19) for all quadrature nodes $\left\{\omega_{\ell,j}^{2c}\right\}$, which can be accomplished efficiently using the Non-equispaced Fast Fourier Transform (NFFT) [3, 8, 28, 5]. The computational complexity of such algorithms is

$$O\left(L^2 \log L + \log \epsilon^{-2} M\right), \tag{35}$$

where L is the sampling rate, ϵ is the required accuracy of the transform, and M is the number of evaluation points, which is equal to the number of quadrature nodes (34). Next, we notice that for every value of ℓ, the inner sum in (25)

$$C_\ell^N \triangleq \sum_{j=1}^{\mathcal{N}_\theta^\ell} \phi_m^c(\omega_{\ell,j}^{2c}) e^{-\imath N\theta} \tag{36}$$

can be computed for multiple values of N using the FFT in $O(\mathcal{N}_\theta^\ell \log \mathcal{N}_\theta^\ell)$ operations, resulting in total of $O\left(\sum_{l=1}^{\mathcal{N}_r} \mathcal{N}_\theta^\ell \log \mathcal{N}_\theta^\ell\right)$ operations for all required values of ℓ. By (30) and (33), this results in a computational complexity of $O(L^2 \log L)$ operations for $c = \pi L$. Lastly, the approximated expansion coefficients (25) are given by the remaining outer sum

$$\tilde{a}_{N,n}^m = \sum_{\ell=1}^{\mathcal{N}_r} \tilde{\mathcal{W}}_\ell^{2c} \frac{R_{N,n}^c(r_\omega^\ell) r_\omega^\ell}{\mathcal{N}_\theta^\ell} C_\ell^N, \tag{37}$$

which can be computed for all indices $\{N, n\} \in \Omega_T$ in $\mathcal{N}_r |\Omega_T|$ operations, where $|\Omega_T|$ denotes the cardinality of the index set Ω_T. From (16) and (33), the computational complexity of the last step is essentially $O(L^3)$, which thus governs the computational complexity of the entire procedure described in this section. We end this discussion with the interesting observation, that the number of operations in the computation of (37) can be further reduced by exploiting the rapid decay of the functions $R_{N,n}^c(r)$ with the radius r (see Figure 1b for an illustration), due to which a substantial number of the terms involving $R_{N,n}^c(r_\omega^\ell)$ can be discarded for certain sets of the radii $\left\{r_\omega^\ell\right\}$ and indices $\{(N, n)\}$. To exemplify this point, for $T = 1$, $L = 150$ and $c = \pi L$, we have that about 30% of the values $\left\{\left|R_{N,n}^c(r_\omega^\ell)\right|\right\}$ are below 10^{-12}, and therefore can be safely discarded from (37).

4 Steerable-PCA procedure

As we have discussed in the Introduction, steerable-PCA extends the classical PCA by artificially including in the analysed dataset all (infinitely many) planar rotations of each image. In the previous sections, we have established a framework for expanding images using PSWFs, under the assumption that the images to be approximated are sufficiently concentrated in space and frequency. As it turns out that PSWFs constitute an excellent basis for expanding images which are localized in space and frequency (see (14), (15) and [13]), we use them in this section to construct optimized basis functions for a given set of images and all of their rotations.

In classical PCA, one has to perform an eigen-decomposition of the covariance matrix of the data in order to obtain the principal components and their associated eigenvalues. The gist of this section is that, by expanding our images using a finite set of steerable basis functions (PSWFs for example), the covariance matrix of the images and all of their rotations transforms into a block-

diagonal matrix in the domain of the expansion coefficients. Thus, we are able to break down our problem into many sub-problems, by computing the eigen-decomposition of the blocks of the diagonal one-by-one. This idea has appeared already in [37] and [36], and is in fact applicable to any steerable orthonormal basis, including the PSWFs basis. Yet, as mentioned in the introduction, replacing the Fourier Bessel basis with the PSWFs basis results in a construction whose properties are fundamentally different. In particular, in the derivation below we consider the principal components as bivariate square integrable functions, instead of considering only their expansion coefficients. This allows us to make the connection between the resulting principal components and the continuous images $I_m(x)$ (prior to sampling), and provide appropriate error bounds.

Recall that $I_m(x)$ stands for the m'th continuous image in our dataset, and let us denote by $I_m^\varphi(x)$ its planar rotation by an angle $\varphi \in [0, 2\pi)$. Ideally, we would like to obtain basis functions that best approximate (in the sense of $\mathcal{L}^2(\mathbf{D})$) the images $I_m^\varphi(x)$ for all rotation angles $\varphi \in [0, 2\pi)$. However, since we do not have access to the underlying continuous images $I_m(x)$, nor their rotations $I_m^\varphi(x)$, we replace them by their PSWFs-based approximations $\hat{I}_m(x)$ (see (10)) and their rotations $\hat{I}_m^\varphi(x)$ respectively. We will now proceed to describe the procedure of obtaining the steerable principal components of the set of approximated images, and at the end of the section make the connection with the original underlying images.

We define the k'th principal component $g_k(x) \in \mathcal{L}^2(\mathbf{D})$ of our set of approximated images $\left\{\hat{I}_m^\varphi \ : \ m = 0, \ldots, M-1, \ \varphi \in [0, 2\pi)\right\}$ as

$$\begin{aligned} g_k(x) = \underset{g(x)}{\operatorname{argmax}} \frac{1}{M} \sum_{m=0}^{M-1} \frac{1}{2\pi} \int_0^{2\pi} \left| \left\langle \left[\hat{I}_m^\varphi(x) - \hat{\mu}(x) \right]^*, g(x) \right\rangle_{\mathcal{L}^2(\mathbf{D})} \right|^2 d\varphi, \\ \text{s.t.} \quad \|g(x)\|_{\mathcal{L}^2(\mathbf{D})} = 1, \quad \langle g(x), g_i(x) \rangle_{\mathcal{L}^2(\mathbf{D})} = 0, \quad 0 \le i < k \end{aligned} \tag{38}$$

where $\langle \cdot, \cdot \rangle_{\mathcal{L}^2(\mathbf{D})}$ denotes the standard inner product on $\mathcal{L}^2(\mathbf{D})$, and $\hat{\mu}(x)$ is the average of all approximated images and their planar rotations

$$\hat{\mu}(x) = \frac{1}{M} \sum_{m=0}^{M-1} \frac{1}{2\pi} \int_0^{2\pi} \hat{I}_m^\varphi(x) d\varphi. \tag{39}$$

In other words, a function $g_k(x)$ is expected to maximize the average projection norm (defined over all images and their rotations), such that $\{g_k(x)\}$ forms a set of orthonormal functions over $\mathcal{L}^2(\mathbf{D})$. The functions $\{g_k(x)\}$ are named the steerable principal components of our dataset. The formulation (38) differs from the classical formulation of PCA (or Karhunen Loeve expansion) only in the additional integration over the rotation angle φ, which has the interpretation of adding all (infinitely many) rotations of all images to our dataset.

Now, for the operation of rotating an approximated image $\hat{I}_m(x)$ by an angle φ, we have

$$\hat{I}_m^{\varphi}(r,\theta) = \hat{I}_m(r,\theta-\varphi) = \sum_{N,n\in\Omega} \hat{a}_{N,n}^m e^{-\imath N\varphi} \hat{\psi}_{N,n}^c(r,\theta), \tag{40}$$

where we have used (10) and the polar separation property of (3). This essentially means that the PSWFs expansion coefficients of $\hat{I_m^{\varphi}}$, which we denote by $\hat{a}_{N,n}^{m,\varphi}$, have the simple form of

$$\hat{a}_{N,n}^{m,\varphi} = \hat{a}_{N,n}^m e^{-\imath N\varphi}. \tag{41}$$

We note that since the coefficients $\hat{a}_{N,n}^m$ can be approximated to high precision by the method described in Section 3 (formula (25)), for ease of notation we do not distinguish between the two. Continuing, as for the evaluation of the mean (approximated) image of (39), using (40) we have

$$\hat{\mu}(x) = \sum_{n=1}^{n_0} \hat{\mu}_{0,n} \hat{\psi}_{0,n}^c(x), \tag{42}$$

where

$$\hat{\mu}_{0,n} = \frac{1}{M} \sum_{m=0}^{M-1} \hat{a}_{0,n}^m, \tag{43}$$

and n_0 in (42) stands for the largest radial index n such that $(0,n) \in \Omega$ (see (8)). Next, let us expand $g_k(x)$ of (38) as a series of PSWFs

$$g_k(x) = \sum_{N,n} \hat{g}_{N,n}^k \hat{\psi}_{N,n}^c(x). \tag{44}$$

It is clear from (38) that the coefficients $\hat{g}_{N,n}^k$ for $(N,n) \notin \Omega$ do not affect that optimization problem. Thus, is sufficient to consider the truncated expansion

$$g_k(x) = \sum_{N,n\in\Omega} \hat{g}_{N,n}^k \hat{\psi}_{N,n}^c(x). \tag{45}$$

Therefore, it follows that we can rewrite (38) as an optimization problem in the domain of the PSWFs expansion coefficients as

$$\begin{aligned} \hat{g}_k &= \underset{\hat{g}}{\operatorname{argmax}}\, \hat{\lambda}_k \\ \text{s.t.} \quad \hat{\lambda}_k &= \frac{1}{M} \sum_{m=0}^{M-1} \frac{1}{2\pi} \int_0^{2\pi} \left| (\hat{a}_m^{\varphi} - \hat{\mu})^* A \hat{g} \right|^2 d\varphi \\ & \|\hat{g}\| = 1, \;\; \hat{g}^* \hat{g}_i = 0, \;\; 0 \le i < k, \end{aligned} \tag{46}$$

where $\hat{a}_m^{\varphi}$, $\hat{\mu}$ and $\hat{g}_k$ are column vectors consisting of the expansion coefficients of $\hat{I}_m^{\varphi}(x)$, $\hat{\mu}(x)$

and $g_k(x)$ respectively, $(\cdot)^*$ is the conjugate-transpose operator, and A is a diagonal matrix whose diagonal is given by

$$A_{(N,n),(N,n)} = \left(\frac{c}{2\pi L}\right)^2 \left|\alpha^c_{N,n}\right|^2, \tag{47}$$

with $\alpha^c_{N,n}$ being the eigenvalue corresponding to the PSWF $\psi^c_{N,n}(x)$ (see (2)). To derive (46), we used (9) and the orthonormality property of the PSWFs on the unit disk. In order to solve the optimization problem (46), we first define the modified coefficients vector $\hat{b}^{\varphi}_m$ as

$$\hat{b}^{\varphi}_m \triangleq A^* \left(\hat{a}^{\varphi}_m - \hat{\mu}\right), \tag{48}$$

which is a column vector of length $|\Omega|$. Consequently, the solution to (46) is given by the eigen-decomposition of the matrix

$$C \triangleq \frac{1}{M} \sum_{m=0}^{M-1} \frac{1}{2\pi} \int_0^{2\pi} \hat{b}^{\varphi}_m \left(\hat{b}^{\varphi}_m\right)^* d\varphi, \tag{49}$$

which is a Hermitian positive semi-definite matrix of size $|\Omega| \times |\Omega|$. Specifically, $\hat{\lambda}_k$ and $\hat{g}_k$ are the eigenvalues and eigenvectors, respectively, of the matrix C, where the eigenvalues $\hat{\lambda}_1 \geq \hat{\lambda}_2 \geq \ldots \geq \hat{\lambda}_{|\Omega|} \geq 0$ are ordered in a non-increasing order. We refer to C as our rotationally-invariant covariance matrix (replacing the standard covariance matrix used in classical PCA). When considering (49) in conjunction with (48) and (41), we get that C has a block-diagonal structure, and exact analytic evaluation of (49) states that

$$C_{(N,n),(N',n')} = \delta_{N,N'} \frac{1}{M} \sum_{m=0}^{M-1} \hat{b}^m_{N,n} \left(\hat{b}^m_{N',n'}\right)^*, \tag{50}$$

where $\hat{b}^m_{N,n}$ denotes the entry of the vector $\hat{b}^0_m$ from (48) which corresponds to the index (N,n). We notice from (50) that we can order the blocks of C according to the index N, such that the matrix C is written as a block-diagonal matrix

$$C = B^{(1)} \oplus B^{(2)} \ldots \oplus B^{(N_\Omega)}, \quad B^{(N)}_{n,n'} = \frac{1}{M} \sum_{m=0}^{M-1} \hat{b}^m_{N,n} \left(\hat{b}^m_{N,n'}\right)^*, \tag{51}$$

where N_Ω is the largest index N such that $(N, n') \in \Omega$ for some radial index n'. For simplicity of the presentation, we provide the reader with the explicit expressions for the coefficients $\hat{b}^m_{N,n}$:

$$\hat{b}^m_{N,n} = \left(\frac{c}{2\pi L}\right)^2 \left|\alpha^c_{N,n}\right|^2 \cdot \begin{cases} \hat{a}^m_{N,n} - \frac{1}{M} \sum_{m'=0}^{M-1} \hat{a}^{m'}_{0,n} & N = 0, \\ \hat{a}^m_{N,n} & N \neq 0. \end{cases} \tag{52}$$

When working with real-valued images, we have that $\hat{a}^m_{N,n} = \left(\hat{a}^m_{-N,n}\right)^*$, and thus it is sufficient

to evaluate the rotationally-invariant covariance matrix $C_{(N,n),(N',n')}$ only for non-negative values of N and N'. Also, since the matrix C enjoys a block-diagonal structure, we are able to obtain its eigen-decomposition through the eigen-decomposition of its blocks $\left\{B^{(N)}\right\}$. Essentially, the eigenvalues and eigenvectors of C can be obtained as the union of the eigenvalues and eigenvectors of the blocks of C while appropriately zero-padding the eigenvectors of the blocks to length $|\Omega|$. Moreover, since the number of chosen basis functions for each angular index N decreases with N (see illustration in [13]), the block sizes decrease as N increases, which reduces the computational complexity of the eigen-decomposition as we process the blocks.

After computing the eigen-decomposition of C, our eigenfunctions $g_k(x)$ can be recovered from their PSWFs expansion coefficients using (45). Specifically, for each pair of eigenvalue and eigenvector $(\hat{\lambda}_k, \hat{g}_k)$ of (51), as $\hat{g}_k$ is nonzero only on some block $B^{(N_k)}$, we can write $g_k(x)$ in polar coordinates as

$$g_k(r,\theta) = \sum_{n=1}^{n_k} \hat{g}^k_{N_k,n} \hat{\psi}^c_{N_k,n}(r,\theta) = \left(\frac{c}{2\pi L}\right) \alpha_{N_k,n} \left(\sum_{n=1}^{n_k} \hat{g}^k_{N_k,n} R^c_{N_k,n}(r)\right) e^{\imath N_k \theta}, \tag{53}$$

where n_k stands for the largest index n such that $(N_k, n) \in \Omega$, and $\hat{g}^k_{N,n}$ are the entries of the eigenvector $\hat{g}_k$ (which are the expansion coefficients of $g_k(x)$ from (45) restricted to $N = N_k$). We note that (53) immediately unveils the fact (mentioned in the introduction) that the resulting principal components are also steerable.

As in classical PCA, the eigenfunctions $g_k(x)$ form the optimal basis for expanding our approximated images $\left\{\hat{I}_m(x)\right\}$ and their rotations, such that if we choose to expand our approximated images using the set of K basis functions $\{g_k\}_{k=1}^K$ as

$$\tilde{I}^\varphi_m(x) = \hat{\mu}(x) + \sum_{k=1}^{K} \left\langle \hat{I}^\varphi_m(x) - \hat{\mu}(x), g_k(x) \right\rangle_{\mathcal{L}^2(\mathbf{D})} g_k(x), \tag{54}$$

then, we have that

$$\frac{1}{M} \sum_{m=0}^{M-1} \frac{1}{2\pi} \int_0^{2\pi} \left\| \hat{I}^\varphi_m(x) - \tilde{I}^\varphi_m(x) \right\|^2_{\mathcal{L}^2(\mathbf{D})} d\varphi \leq \sum_{k=K+1}^{|\Omega|} \hat{\lambda}_k, \tag{55}$$

where $\hat{\lambda}_k$ is the k'th eigenvalue of C. Also, (53) shows that the expansion coefficients in (54), given by $\left\langle \hat{I}^\varphi_m(x) - \hat{\mu}(x), g_k(x) \right\rangle_{\mathcal{L}^2(\mathbf{D})}$, can be computed efficiently by

$$\left\langle \hat{I}^\varphi_m(x) - \hat{\mu}(x), g_k(x) \right\rangle_{\mathcal{L}^2(\mathbf{D})} = \sum_{n=1}^{n_k} \hat{b}^m_{N_k,n} \left(\hat{g}^k_{N_k,n}\right)^*. \tag{56}$$

Now, we point out that the steerable principal components were computed for the set of approxi-

mated images $\left\{\hat{I}_m(x)\right\}$, and not the set of the underlying images $\{I_m(x)\}$, which suggests that (54) is accurate (in the sense of (55)) only for the set of the approximated images. Nonetheless, the PSWFs approximation scheme from Section 2 provides us with strict error bounds when expanding images localized in space and frequency. Therefore, when using (54) to approximate our images $I_m(x)$, the error norm (averaged over all images and rotations) can be bounded by joining (55) and (15). In particular, for a truncation parameter T, we have that

$$\frac{1}{M}\sum_{m=0}^{M-1}\frac{1}{2\pi}\int_0^{2\pi}\left\|I_m^{\varphi}(x)-\tilde{I}_m^{\varphi}(x)\right\|_{\mathcal{L}^2(\mathbf{D})}^2 d\varphi \leq \sum_{k=K+1}^{|\Omega_T|}\hat{\lambda}_k + 2\mathcal{E}(\varepsilon,\delta_c,T)\sqrt{\sum_{k=K+1}^{|\Omega_T|}\hat{\lambda}_k}+\mathcal{E}^2(\varepsilon,\delta_c,T), \quad (57)$$

where $\tilde{I}_m^{\varphi}(x)$ is the expansion via the steerable principal components (54), and $\mathcal{E}(\varepsilon,\delta_c,T)$ is the approximation error term of (15) for the truncated series of PSWFs. In essence, (57) asserts that if the images are sufficiently localized in space and frequency, then for an appropriate truncation parameter T, the error in expanding $I_m(x)$ using the steerable principal components computed from $\hat{I}_m(x)$ is close to the smallest possible error given by (55).

5 Algorithm summary and computational cost

We summarize the algorithms for evaluating the expansion coefficients $\hat{a}_{N,n}^m$ and $\tilde{a}_{N,n}^m$ (corresponding to the direct method from Section 2 and the efficient method of Section 3, respectively), and the steerable-PCA procedure described in Section 4, in Algorithms 1, 2, and 3, respectively.

Algorithm 1 Evaluating PSWFs expansion coefficients (direct method)

1: **Required:** M images $\left\{I_m(\frac{k}{L})\right\}_{m=1}^M$ sampled on a Cartesian grid of size $(2L+1)\times(2L+1)$ with $\frac{k}{L}\in\mathbf{Q}$, where $\mathbf{Q}=[-1,1]\times[-1,1]$.
2: **Precomputation:**
   1. Choose a bandlimit c ($\leq \pi L$) and a truncation parameter T.
   2. Evaluate the PSWFs $\psi_{N,n}^c(\frac{k}{L})$ and their eigenvalues $\alpha_{N,n}^c$ according to [30] for $(N,n)\in\Omega_T$ (see (12)), where $\frac{k}{L}\in\mathbf{D}$.
   3. Compute the modified PSWFs $\hat{\psi}_{N,n}^c(\frac{k}{L})=\left(\frac{c}{2\pi L}\right)\alpha_{N,n}^c\psi_{N,n}^c(\frac{k}{L})$, $(N,n)\in\Omega_T$.
3: For every $m=0,\ldots,M-1$ and $(N,n)\in\Omega_T$, compute the expansion coefficients $\hat{a}_{N,n}^m=\sum_{\frac{k}{L}\in\mathbf{D}} I_m(\frac{k}{L})\left[\hat{\psi}_{N,n}^c(\frac{k}{L})\right]^*$.

We now turn our attention to the computational complexity of Algorithms 1, 2, and 3. We omit the pre-computation steps from the complexity analysis as they can be performed only once per setup, and do not depend on the specific images.

Since we have $O(L^2)$ PSWFs in our expansions (corresponding to the indices in the set Ω_T) and

Algorithm 2 Evaluating PSWFs expansion coefficients (efficient method)

1: **Required:** M images $\left\{I_m(\frac{k}{L})\right\}_{m=1}^{M}$ sampled on a Cartesian grid of size $(2L+1)\times(2L+1)$ with $\frac{k}{L}\in\mathbf{Q}$, where $\mathbf{Q}=[-1,1]\times[-1,1]$.

2: **Precomputation:**

1. Choose a bandlimit c ($\leq \pi L$) and a truncation parameter T.
2. Choose the number of radial nodes $\mathcal{N}_r$ and angular nodes $\mathcal{N}_\theta^\ell$ for $1\leq\ell\leq\mathcal{N}_r$ (according to (28) and (29) or similar relaxed conditions).
3. Compute the quadrature nodes $\omega_{\ell,j}^{2c}=\left(r_\omega^\ell,\theta_\omega^{\ell,j}\right)$ and weights $\mathcal{W}_{\ell,j}^{2c}$ for $\ell=1,\ldots,\mathcal{N}_r$ and $j=1,\ldots,\mathcal{N}_\theta^\ell$, as described in [30].
4. Evaluate the radial part of the PSWFs $R_{N,n}^c(r)$ at the radial quadrature nodes r_ω^ℓ for $(N,n)\in\Omega_T$ (see (12)).

3: For $m=0,\ldots,M-1$, compute $\phi_m^c(\omega_{\ell,j}^{2c})$ from (19) by NFFT.

4: For $m=0,\ldots,M-1$ and $(N,n)\in\Omega_T$, compute the expansion coefficients $\tilde{a}_{N,n}^m$ via equations (36) and (37).

Algorithm 3 PSWFs-based steerable-PCA

1: **Required:** PSWFs expansion coefficients $\left\{\hat{a}_{N,n}^m\right\}_{m=0}^{M-1}$ for $(N,n)\in\Omega_T$ (see (12)) and a bandlimit c.

2: **Precomputation:**

1. Evaluate the PSWFs $\psi_{N,n}^c(\frac{k}{L})$, $\frac{k}{L}\in\mathbf{D}$, and their eigenvalues $\alpha_{N,n}^c$ for $(N,n)\in\Omega_T$ according to [30].
2. Compute the modified PSWFs $\hat{\psi}_{N,n}^c(\frac{k}{L})=\left(\frac{c}{2\pi L}\right)\alpha_{N,n}^c\psi_{N,n}^c(\frac{k}{L})$.

3: Compute the expansion coefficients of the mean image $\hat{\mu}_{0,n}=\frac{1}{M}\sum_{m=0}^{M-1}\hat{a}_{0,n}$ for $n=1,\ldots,n_0$ where n_0 is the largest index n such that $(0,n)\in\Omega_T$.

4: Update $\hat{a}_{0,n}^m\leftarrow\hat{a}_{0,n}^m-\hat{\mu}_{0,n}$.

5: Compute the coefficients $\hat{b}_{N,n}^m=\left(\frac{c}{2\pi L}\right)^2\left|\alpha_{N,n}^c\right|^2\hat{a}_{N,n}^m$.

6: Compute the eigenvalues $\hat{\lambda}_1,\ldots,\hat{\lambda}_{|\Omega_T|}$ and eigenvectors $\hat{g}_1,\ldots,\hat{g}_{|\Omega_T|}$ of the matrix C (from (50)) by diagonalizing each of its blocks separately.

7: Compute the sampled basis functions $g_\ell(\frac{k}{L})=\sum_{N,n\in\Omega_T}\hat{g}_{N,n}^\ell\hat{\psi}_{N,n}^c(\frac{k}{L})$, where $\hat{g}_{N,n}^\ell$ denotes the entry of $\hat{g}_\ell$ corresponding to the pair (N,n).

8: The expansion coefficients of $I_m(x)$ in the steerable-PCA basis are given by $d_k^m=\sum_{N,n\in\Omega_T}\hat{a}_{N,n}^m(\hat{g}_{N,n}^k)^*$.

$O(L^2)$ equally-spaced Cartesian samples in the unit disk, computing the expansion coefficients of M images using the direct approach of Algorithm 1 requires $O(ML^4)$ operations. On the other hand, Algorithm 2 allows us to obtain the expansion coefficients in $O(ML^3)$ operations, since the NFFT in step 3 requires $O(L^2 \log L)$ operations and step 4 can be implemented using $O(L^2 \log L + L^3)$ operations, for each image.

Although both our method and the method described in [36] (based on Fourier-Bessel basis functions) have the same order of computational complexity, our method enjoys a twofold asymptotic speedup in the coefficients' evaluation, and often it runs about three/four times faster. The twofold asymptotic speedup is because we need asymptotically only half the number of radial nodes for the numerical integration (see analysis in Section 3 and Appendix C) compared to the Gaussian quadratures used in [36], which dictates the constant in the leading term of the computational complexity. The greater speedup observed in practice stems from the fact that practically, the most time-consuming operation in both methods is the NFFT, which depends heavily on the total number of quadrature nodes. In our integration scheme, the total number of quadrature nodes is about 1/4 of the number of nodes in [36].

Next, as for the computational complexity of the steerable-PCA (Algorithm 3), forming the blocks of the matrix C requires $O(ML^3)$ operations, and then the eigen-decomposition of each block requires $O(L^3)$ operations, where there are $O(L)$ different blocks in the matrix. Therefore, obtaining the eigenvalues and eigenvectors in step 6 of Algorithm 3 requires $O(ML^3 + L^4)$ operations. As pointed out in [37] and [36], the eigenvalues and eigenvectors can be also obtained from the SVD of the matrices of coefficients $\hat{b}^m_{N,n}$ corresponding to each block of C.

Following (53), if we have $O(L^2)$ basis functions $g_k(x)$, their evaluation on the Cartesian grid requires $O(L^5)$ operations. Sometimes it may be more convenient to evaluate these basis functions on a polar grid instead, in which case the computational complexity reduces to $O(L^4)$.

Lastly, computing the expansion coefficients of the images in the steerable basis via (56) requires $O(ML^3)$ operations, since $O(L)$ operations are required to compute a single expansion coefficient for every image.

We point out though, that often only a small fraction of the basis functions are chosen for subsequent processing (via their eigenvalues), so the contribution of steps 7 and 8 in Algorithm 3 to the overall running time of the steerable-PCA procedure is usually negligible.

To summarize, the computational complexity of the entire procedure (computing PSWFs expansion coefficients + steerable-PCA), when using Algorithm 2 to compute the expansion coefficients, is $O(ML^3 + L^5)$ for sampling the steerable principal components on the Cartesian grid, and $O(ML^3 + L^4)$ when sampling them on a polar grid. It is important to note that, although Algorithm 1 for evaluating PSWFs expansion coefficients provides us with an inferior order of computational complexity (compared to Algorithm 2), it is simpler to implement and may still run faster (due to optimized implementations of the scalar product on CPUs and GPUs), particularly for small values of L.

6 Steerable-PCA in the presence of noise

Up to this point, we have presented a method for computing the steerable-PCA of a set of images localized in space and frequency, sampled on a Cartesian grid. In many practical settings however, the images are corrupted with noise. Therefore, it is beneficial to understand the impact this noise has on the PSWFs expansion coefficients, and in this aspect, it is generally convenient if the transformation to the expansion coefficients does not alter the spectrum of the noise. In addition, it is also favourable if the basis functions obtained from steerable-PCA form an orthonormal basis when sampled on the Cartesian grid, as would be if we instead applied the standard PCA on the Cartesian samples of the images. In this section, we demonstrate numerically that for sufficiently localized images in space/frequency, for which we can choose a truncation parameter $T \gg 1$ (see (12)), the transformation to the PSWFs expansion coefficients is essentially orthonormal, and the set of steerable principal components sampled on the Cartesian grid is also practically orthonormal. In particular, the higher the truncation parameter T, the closer is our steerable-PCA to orthonormality.

Let us denote by I_m a column vector of the clean (Cartesian) samples of the m'th image. Now, suppose that our images are corrupted by zero mean additive noise, such that

$$\tilde{I}_m = I_m + \xi_m, \tag{58}$$

where ξ_m denotes the column vector of the noise samples added to the m'th image. From (11), we compute our approximated expansion coefficients by

$$\hat{a}_m = \left(\hat{\psi}^c\right)^* \tilde{I}_m = \left(\hat{\psi}^c\right)^* I_m + \left(\hat{\psi}^c\right)^* \xi_m, \tag{59}$$

where the operator $(\cdot)^*$ stands for the conjugate-transpose, and $\hat{\psi}^c$ denotes a matrix whose columns contain samples of the modified PSWFs (defined in (9)) inside the unit disk, with different columns corresponding to different pairs of indices (N, n) chosen via the truncation rule of (12). Let us define the additive noise vector in the expansion coefficients of the m'th image by

$$\hat{\xi}_m = \left(\hat{\psi}^c\right)^* \xi_m. \tag{60}$$

Since the transformation from the sampled images to the expansion coefficients is simply a multiplication by the conjugate transpose of the matrix $\hat{\psi}^c$, it follows that the covariance matrix of the additive noise in the expansion coefficients is given by

$$R_{\hat{\xi}} = \left(\hat{\psi}^c\right)^* R_\xi \hat{\psi}^c, \tag{61}$$

where $\hat{\xi}$ denotes the random vector of additive noise in the expansion coefficients, and ξ denotes

the random vector of additive noise in the images, whose covariance matrix is R_ξ. We note that by assuming that the noise distribution is the same for all images, we drop the 'm' index from the notation of the random noise vectors. Now, if the additive noise in (58) is white ($R_\xi = \sigma^2 I$), and in order to preserve the spectrum of the noise, we would require the matrix

$$H_c \triangleq \left(\hat{\psi}^c\right)^* \hat{\psi}^c \tag{62}$$

to be as close as possible to the identity matrix, or equivalently, that the eigenvalues of the matrix H_c are as close as possible to 1.

Moreover, we would like the sampled steerable principal components, computed in step 7 of Algorithm 3, to be orthogonal. Let us recall that the k'th basis function obtained from the steerable-PCA procedure and sampled on the Cartesian grid inside the unit disk (see step 7 in Algorithm 3), is given by

$$g_k = \hat{\psi}^c \hat{g}_k, \tag{63}$$

where $\hat{g}_k$ is the column vector of PSWFs expansion coefficients for the principal component $g_k(x)$, and g_k is a column vector of the equally-spaced samples of $g_k(x)$ inside the unit disk. It is clear that since the vectors $\{\hat{g}_k\}$ were obtained from the steerable-PCA procedure, they form an orthonormal set (see (46)), and so the orthonormality of the set of sampled principal components $\{g_k\}$ directly depends on the spectrum of H_c.

As of the matrix H_c, from its definition it is Hermitian and positive semi-definite, and therefore it has non-negative real-valued eigenvalues $\nu_1, \nu_2, \ldots, \nu_{\Omega_T}$. To determine how close H_c is to the identity matrix, we evaluated numerically the maximal distance (in absolute value) between the eigenvalues $\{\nu_k\}$ and 1, and the results are plotted in Figure 6 for various values of L, T and the bandlimit c. It is noteworthy that for $T = 10^6$ the eigenvalues of H_c differ from 1 by about 10^{-12}. We have also observed that the spectrum of the matrix H_c is related to the truncation parameter T by

$$\max_k |1 - \nu_k| \approx T^{-2}, \tag{64}$$

for $T \geq 1$. This observation is exemplified numerically in Figure 5. Having established empirically that the transform H_c in (62) is essentially orthogonal for sufficiently large values of T, it is important to mention that if the images are sufficiently localized in space and frequency, it is often possible to increase the value of T such that the transform is sufficiently orthogonal without significantly degrading the quality of approximation.

7 Numerical experiments

We begin by conducting an experiment to demonstrate the running times of our algorithm for large datasets. We generated $20,000$ white noise images of varying sizes, and compared the running time

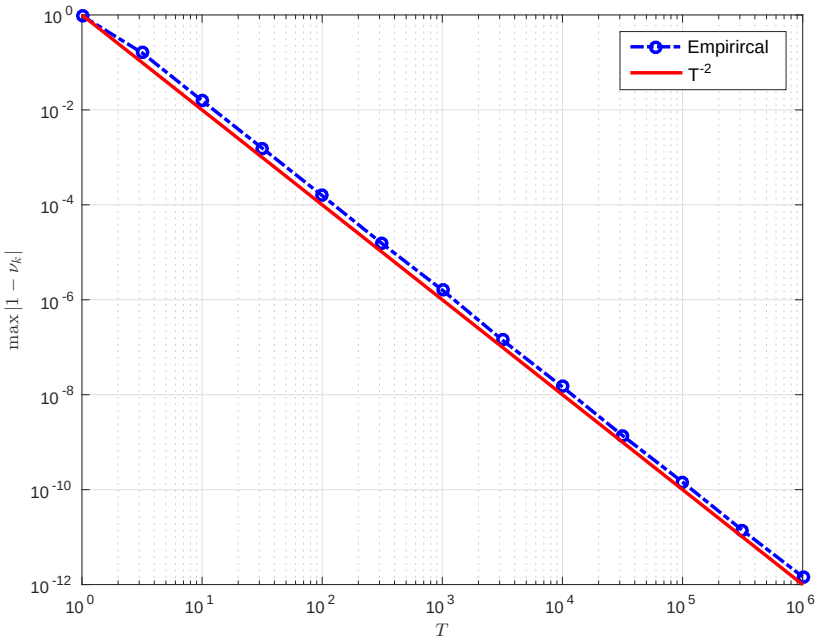

Figure 5: Measured deviation of the eigenvalues of H_c from 1 as a function of T, versus the function T^{-2}, for $L = 96$ and $c = \pi L$. We obsereve that the truncation parameter T gives us direct control over the spectrum of H_c with $\max_k |1 - \nu_k| \approx T^{-2}$, for $T \geq 1$.

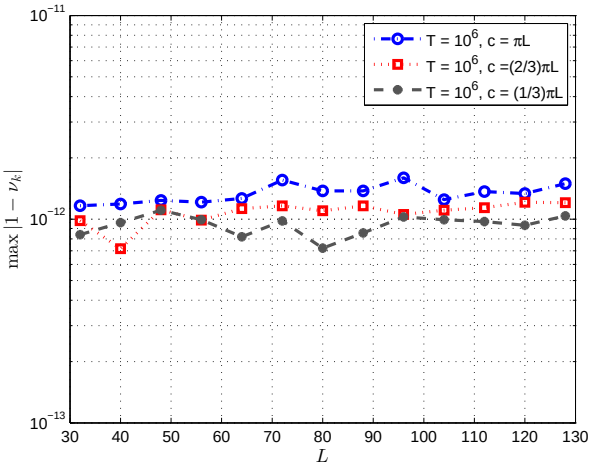

(a) $T = 10^6$

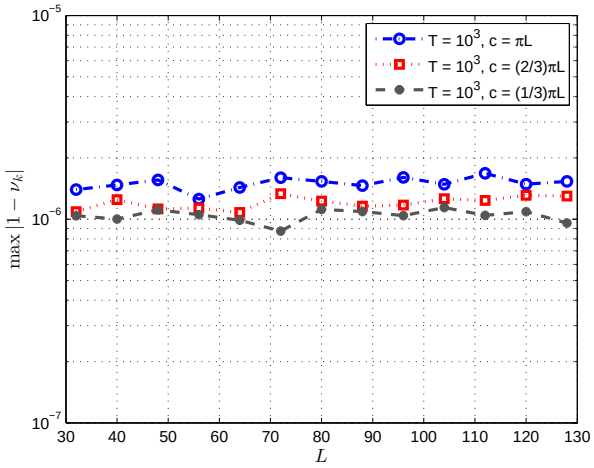

(b) $T = 10^3$

Figure 6: Measured deviation of the eigenvalues of H_c from 1 for different values of the truncation parameter T, the bandlimit c, and the sampling resolution L. We notice that the deviation from orthogonality remains approximately constant for different values of L and c. Specifically, for $T = 10^6$ we notice that H_c is practically orthogonal.

L	PSWFs-direct (Algs. 1+3)	PSWFs-efficient (Algs. 2+3)	FFBsPCA [36]
32	6	25	47
64	95	59	151
96	491	117	363
128	1625	204	697

Table 1: Comparison of algorithms' running times, for $20,000$ images of size $2L \times 2L$ consisting of white noise, for $T = 10$ and several values of L. All timings are given in seconds.

of our algorithm with the FFBsPCA algorithm of [36], where the coefficients' evaluation for our algorithm was carried out using both Algorithm 1 (direct method) and Algorithm 2 (efficient method). All of the algorithms were implemented in Matlab, and were executed on a dual Intel Xeon X5560 CPU (8 cores in total), with 96GB of RAM running Linux. Whenever possible, all 8 cores were used simultaneously, either explicitly using Matlab's `parfor`, or implicitly, by employing Matlab's implementation of BLAS, which takes advantage of multi-core computing. As for the NFFT implementation, we used the software package [12], with an oversampling of 2 and a truncation parameter $m = 6$ (which provides accuracy close to machine precision). The running times (in seconds) are shown in Table 1. As anticipated, Algorithm 1 runs faster than Algorithm 2 for small image sizes, but becomes significantly slower for larger values of L. As for the running time of our algorithm versus FFBsPCA [36], we have mentioned in Section 5 that our algorithm is asymptotically two times faster, since the number of radial nodes in our quadrature scheme is asymptotically half that of FFBsPCA. In addition, the total number of quadrature nodes in our scheme is about four times less than that of FFBsPCA, and since the NFFT procedure for evaluating the Fourier transform of the sampled images on a polar grid (see (19)) is major time consuming step of both algorithms, it is expected that our algorithm will be faster than FFBsPCA by a factor between 2 and 4, which indeed agrees with the results in Table 1. In all scenarios tested, the vast majority of computation time was spent on obtaining the PSWFs expansion coefficients (Algorithm 2), and only a small fraction of the time on the eigen-decomposition of the rotationally-invariant covariance matrix, even though both parts of the algorithm have the same asymptotic computational complexity. Table 2 summarizes the time spent on the evaluation of the PSWFs expansion coefficients (using Algorithm 2) versus time spent on the eigen-decomposition of the rotationally-invariant covariance matrix .

Next, we demonstrate our algorithms on simulated cryo-electron microscopy (cryo-EM) projection images. In cryo-EM, one is interested in reconstructing a three-dimensional model of a macromolecule (such as a protein) from its two-dimensional projection images taken by an electron microscope. The procedure begins by embedding many such macromolecules in a thin layer of ice (hence the 'cryo' in the name of the procedure), where due to the experimental setup the macromolecules are frozen at random unknown orientations. Then, an electron microscope ac-

L	PSWFs coefficients evaluation (Alg. 2)	Eigen-decomposition (Alg. 3 steps 3-6)
32	24.5	0.5
64	56.5	2.5
96	111	6
128	193	11

Table 2: Running times of coefficients' evaluation and eigen-decomposition for the efficient PSWFs-based method. All timings are given in seconds.

quires two-dimensional projection images of the electron densities of these macromolecules. This procedure can be modelled mathematically as the Radon transform of a volume function evaluated at random viewing directions. Due to the properties of the cryo-EM imaging procedure, each projection image generated by the electron microscope undergoes a convolution with a kernel, referred to as the "contrast transfer function" (CTF), which is known to have a Gaussian envelope [6]. Since the unknown volume is essentially limited in space, and since the behaviour of the CTF dictates that all images are localized in Fourier domain, we conclude that projection images obtained by a cryo-EM procedure are essentially limited in both space and frequency to sufficiently large circular domains. This fact makes these images suitable in exemplifying our steerable-PCA algorithm. In the following experiment, we demonstrate the accuracy of our steerable-PCA algorithm. To this end, we simulated $10,000$ clean projection images from a noiseless three-dimensional density map (EMD-5578 from The Electron Microscopy Data Bank (EMDB) [19]) using the ASPIRE package [1], and obtained steerable principal components using both our fast algorithm and FFBsPCA [36]. Then, we used different numbers of principal components to reconstruct the images, and compared the theoretical error predicted by the residual eigenvalues (see (55)) with the empirical error obtained when comparing the reconstructions to the original images. Obviously, the difference between the two errors is due to the error incurred by the images' expansion scheme (PSWFs or Fourier-Bessel functions). Typical projection images from this dataset can be seen in Figure 7. To simplify the setting and the exposition, we used the projection images as they were obtained from the given volume, and we did not crop, filter, or process them beforehand. Therefore, we assumed a bandlimit $c = \pi L$ throughout the experiment. We note that in general it is possible to crop and filter the images (i.e. choose a smaller bandlimit c) without significantly degrading the quality of the images (up to some prescribed accuracy), thus reducing the computational burden of the algorithms. This can be accomplished either by power density estimation (as demonstrated in [36]), or by employing more sophisticated and dataset-specific estimation techniques for the B-factor (which governs the Gaussian envelope decay of the CTF), see for example [23]. Figure 8 depicts the first 12 eigenfunctions obtained by our steerable-PCA algorithm for this dataset. In Figure 9 we show the relative error norms for the FFBsPCA algorithm and our PSWFs-based algorithm, using several values of T and with different numbers of eigenfunctions in

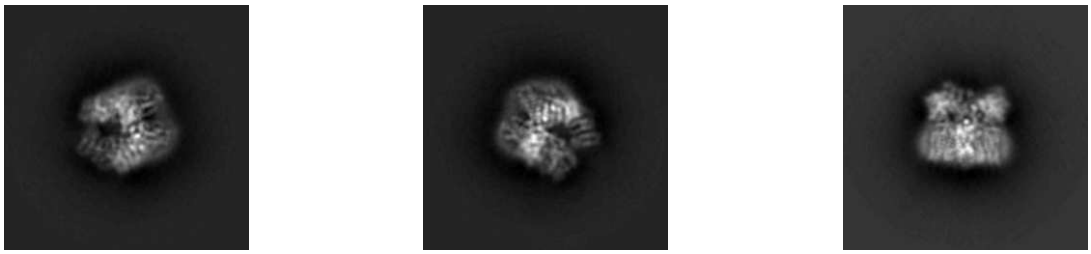

(a) (b) (c)

Figure 7: Sample of 3 simulated noiseless projection images at a resolution of $L = 127$ pixels.

the reconstruction. As expected, the theoretical and empirical relative error norms coincide when a small number of principal components is used in the reconstruction, yet when a large number of principal components is used, the error incurred by the initial series expansion (either using PSWFs or Fourier-Bessel) comes into play and dominates the overall error. As guaranteed by (57) and (15) for space/frequency localized images, smaller values of T lead to smaller approximation errors, and we notice that our PSWFs-based method outperforms the FFBsPCA algorithm in terms of accuracy for $T = 10$ and $T = 10^{-3}$. It is also important to mention that the number of PSWFs taking part in the approximation of each image does not increase significantly when lowering T (see (16)). On the other hand, if one allows the approximation error to be of the order of 10^{-3} to 10^{-4}, then it is also possible to use a very large truncation parameter, such as $T = 10^5$, for which we enjoy superior orthogonality properties of the transform (see Figure 6), as well as shorter PSWFs expansions. Often, such scenarios arise when handling noisy datasets. In a way, the truncation parameter provides us with flexibility to adapt the approximation scheme to our specific setting and needs.

8 Summary and discussion

In this paper, we utilized PSWFs-based computational tools to construct fast and accurate algorithms for obtaining steerable principal components of large image datasets. The accuracy of the algorithms are guaranteed under the assumptions of localization of the images in space and frequency, which are natural assumptions for many datasets, particularly in the field of tomography. For M images, each sampled on a Cartesian grid of size $(2L+1)\times(2L+1)$, the computational complexity for obtaining the steerable principal components is $O(ML^3)$ operations, and their accuracy is of the order of the localization of the images in space and frequency, i.e. the norm of the images outside the unit disk in space, and the norm of their Fourier transform outside a disk of radius c, where c is the chosen bandlimit. We have compared our method with the FFBsPCA algorithm [36], which is considered state-of-the-art for performing steerable-PCA on cryo-EM projection images,

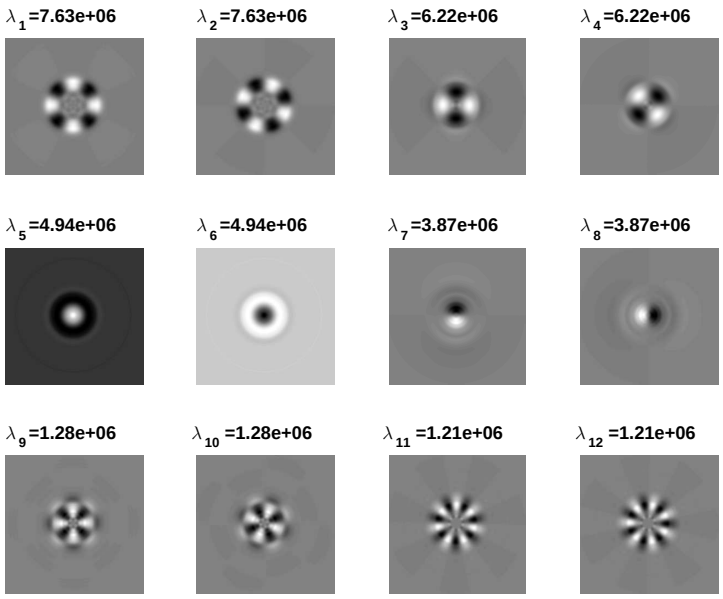

Figure 8: The first 12 eigenfunctions with largest eigenvalues, obtained for $T = 10^{-3}$.

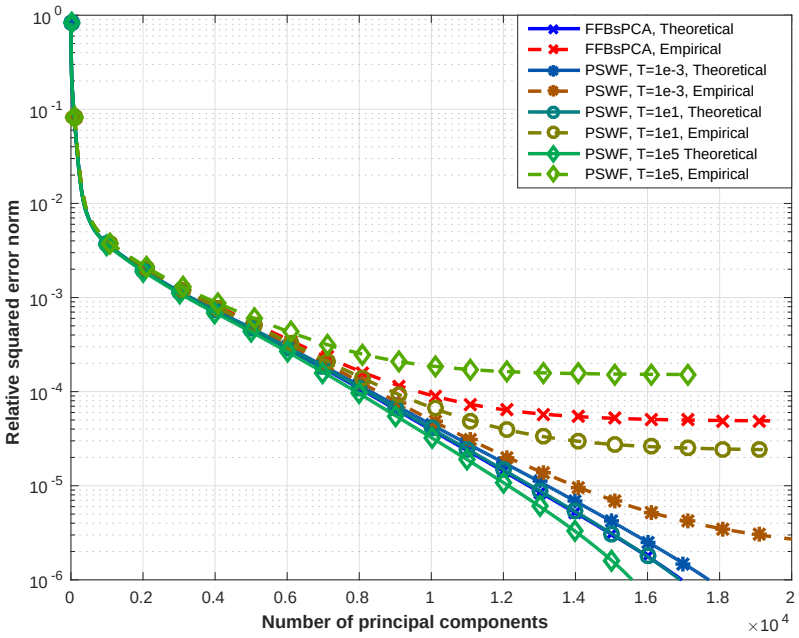

Figure 9: The ratio between the squared error norm and the squared norm of the images, using the FFBsPCA algorithm of [36] and our PSWFs-based algorithm for several values of T. Curves which correspond to theoretical errors are obtained from the residual eigenvalues of the steerable-PCA procedure (see (55)), and empirical errors correspond to measured errors between the reconstructed and original projection images.

and have shown analytically and empirically that our method is both faster and more accurate (for sufficiently small values of T), where's in contrast, the FFBsPCA algorithm provides no analytic guarantees on it's accuracy. We mention that classical operations of windowing and filtering can be subsumed in our procedure to ensure that the images fulfil the requirements of space-frequency localization.

As image resolutions get higher, investigating more efficient methods for processing the images is an important ongoing research task. As the running times of our algorithms are mostly dominated by the task of computing the PSWFs expansion coefficients, reducing the computational complexity of this step from $O(L^3)$ to $O(L^2 \log L)$ for each image, resulting in the same complexity as the 2D FFT, will be a significant progress. Since the radial part of the PSWFs is evaluated to a prescribed accuracy using a finite series of special polynomials which admit a recurrence relation (see [30]), and since the expansion coefficients in terms of these polynomials are obtained from the eigenvectors of a tridiagonal matrix, it is possible to employ the methods described in [33] to derive an $O(L^2 \log L)$ algorithm for computing the PSWFs expansion coefficients of a function. We leave such possible improvements to future work.

Acknowledgements

This research was supported by THE ISRAEL SCIENCE FOUNDATION grant No. 578/14, and by Award Number R01GM090200 from the NIGMS.

Appendix A Simplified PSWFs expansion error bound

Let us define

$$\bar{I}_m(x) \triangleq \begin{cases} I_m(x) & x \in \mathbf{D}, \\ 0 & x \notin \mathbf{D}, \end{cases} \tag{65}$$

i.e. the restrictions of the images $I_m(x)$ to the unit disk. Consider their Fourier transforms

$$\bar{I}_m^{\mathcal{F}}(\omega) \triangleq \mathcal{F}\left\{\bar{I}_m\right\}(\omega) = \int_{\mathbb{R}^2} \bar{I}_m(x) e^{\imath \omega x} dx. \tag{66}$$

Since the Fourier transform is a linear operator, we can write

$$\bar{I}_m^{\mathcal{F}}(\omega) = \mathcal{F}\left\{I_m - \left(I_m - \bar{I}_m\right)\right\}(\omega) = I_m^{\mathcal{F}}(\omega) + \mathcal{F}\left\{I_m - \bar{I}_m\right\}(\omega), \tag{67}$$

where we denoted by $I_m^{\mathcal{F}}(\omega)$ the Fourier transform of $I_m(x)$. Now, it is clear that

$$\left\|\bar{I}_m^{\mathcal{F}}(\omega)\right\|_{\mathcal{L}^2(\omega\notin c\mathbf{D})} \leq \left\|I_m^{\mathcal{F}}(\omega)\right\|_{\mathcal{L}^2(\omega\notin c\mathbf{D})} + \left\|\mathcal{F}\left\{I_m - \bar{I}_m\right\}(\omega)\right\|_{\mathcal{L}^2(\omega\notin c\mathbf{D})} \tag{68}$$

$$\leq \left\|I_m^{\mathcal{F}}(\omega)\right\|_{\mathcal{L}^2(\omega\notin c\mathbf{D})} + \left\|\mathcal{F}\left\{I_m - \bar{I}_m\right\}(\omega)\right\|_{\mathcal{L}^2(\mathbb{R}^2)}. \tag{69}$$

Next, by employing Parseval's identity we obtain that

$$\left\|\mathcal{F}\left\{I_m - \bar{I}_m\right\}(\omega)\right\|_{\mathcal{L}^2(\mathbb{R}^2)} = 2\pi \left\|I_m(x) - \bar{I}_m(x)\right\|_{\mathcal{L}^2(\mathbb{R}^2)}, \tag{70}$$

and thus, by substituting our space/frequency concentration assumptions on the images $I_m(x)$, we have that

$$\left\|\bar{I}_m^{\mathcal{F}}(\omega)\right\|_{\mathcal{L}^2(\omega\notin c\mathbf{D})} \leq \delta_c + 2\pi\varepsilon. \tag{71}$$

It follows that by using our definitions of concentration in space/frequency, we can say that $\bar{I}_m(x)$ is $(1,\bar{\varepsilon})$-concentrated in space, and its Fourier transform is $(c,\bar{\delta}_c)$-concentrated, where

$$\bar{\varepsilon} = 0, \quad \bar{\delta}_c = \delta_c + 2\pi\varepsilon. \tag{72}$$

Finally, since $I_m(x)$ and $\bar{I}_m(x)$ coincide on $\mathbf{D}$, by substituting $\bar{\varepsilon}$ and $\bar{\delta}_c$ into (14) instead of ε and δ_c respectively, we have

$$\left\|I_m(x) - \sum_{N,n\in\Omega_T} \hat{a}_{N,n}^m \hat{\psi}_{N,n}^c(x)\right\|_{\mathcal{L}^2(\mathbf{D})} = \left\|\bar{I}_m(x) - \sum_{N,n\in\Omega_T} \hat{a}_{N,n}^m \hat{\psi}_{N,n}^c(x)\right\|_{\mathcal{L}^2(\mathbf{D})} \tag{73}$$

$$\leq \left(\varepsilon + \frac{\delta_c}{2\pi}\right)(T+4). \tag{74}$$

Appendix B A bound for $\mathcal{N}_\theta^\ell$

Recall that we choose the number of angular nodes per radius $\mathcal{N}_\theta^\ell$ such that it satisfies

$$\sum_{|j|>\mathcal{N}_\theta^\ell} \left|J_j(2cr_\omega^\ell\rho)\right| < C_1\vartheta_q, \tag{75}$$

for every $\rho\in[0,1]$. Using a bound for the Bessel functions [2] together with the fact that $|\rho|\leq 1$, we get

$$\left|J_j(2cr_\omega^\ell\rho)\right| \leq \left(\frac{2cr_\omega^\ell\rho}{2}\right)^j \frac{1}{\Gamma(j+1)} \leq \left(cr_\omega^\ell\right)^j \frac{1}{(j+1)!}, \tag{76}$$

and by using Stirling's approximation [2] (for $n>1$)

$$n! \geq \sqrt{2\pi n}\left(\frac{n}{e}\right)^n > e\left(\frac{n}{e}\right)^n \tag{77}$$

we obtain (for $j > 0$)

$$\left|J_j(2cr_\omega^\ell\rho)\right| \leq \frac{1}{cr_\omega^\ell e}\left(\frac{cr_\omega^\ell e}{j+1}\right)^{j+1}. \tag{78}$$

Then, using the fact that the Bessel functions of the first kind satisfy

$$J_{-n}(x) = (-1)^n J_n(x), \tag{79}$$

we can write

$$\sum_{|j|>\mathcal{N}_\theta^\ell}\left|J_j(2cr_\omega^\ell\rho)\right| = 2\sum_{j>\mathcal{N}_\theta^\ell}\left|J_j(2cr_\omega^\ell\rho)\right| \leq \frac{2}{cr_\omega^\ell e}\sum_{j>\mathcal{N}_\theta^\ell}\left(\frac{cr_\omega^\ell e}{j+1}\right)^{j+1}. \tag{80}$$

We define

$$\gamma_\ell \triangleq cr_\omega^\ell e \tag{81}$$

and choose the number of quadrature nodes per radius as

$$\mathcal{N}_\theta^\ell = \lceil\gamma_\ell\rceil + d, \tag{82}$$

where d is some positive integer and $\lceil\cdot\rceil$ denotes the rounding up operation. It can be easily verified that

$$\frac{\gamma_\ell}{j+1} \leq \frac{\gamma_\ell}{\gamma_\ell + d} \tag{83}$$

whenever $j > \mathcal{N}_\theta^\ell$. Using (80) and (83) we get that

$$\sum_{|j|>\mathcal{N}_\theta^\ell}\left|J_j(2cr_\omega^\ell\rho)\right| \leq \frac{2}{\gamma_\ell}\sum_{j>\mathcal{N}_\theta^\ell}\left(\frac{\gamma_\ell}{\gamma_\ell+d}\right)^{j+1} \leq 2\left(1+\frac{d}{\gamma_\ell}\right)^{-\mathcal{N}_\theta^\ell} \leq 2\left(1+\frac{d}{\gamma_\ell}\right)^{-\gamma_\ell-d} \leq 2e^{-d} \tag{84}$$

where we have used the inequality

$$\left(1+\frac{a}{b}\right)^{-b} < e^{-\frac{ab}{a+b}}, \tag{85}$$

from [2], with $a = (\gamma_\ell + d)\frac{d}{\gamma_\ell}$ and $b = \gamma_\ell + d$. Finally, we can see that in order to satisfy (75) it is sufficient to choose

$$d = d(\vartheta_q) \geq \log\vartheta_q^{-1} + \log\frac{2}{C_1}, \tag{86}$$

and thus

$$\mathcal{N}_\theta^\ell \geq cr_\omega^\ell e + \log\vartheta_q^{-1} + \log\frac{2}{C_1} + 1. \tag{87}$$

Appendix C Behavior and decay properties of $\lambda_{0,k}^{2c}$

Let us first discuss some known results on the behavior of the eigenvalues of the PSWFs in the one-dimensional setting, where they have been thoroughly investigated (see [24] and references therein). The most well-known characterization of these eigenvalues is that they can be divided into three distinct regions of behavior (as a function of their index n) - a flat region, where the (normalized) eigenvalues are essentially 1, a transitional region, where they decay from values close to 1 to values close to 0, and a super-exponential decay region, where they are very close to 0 and decay as $\sim e^{-n\log n}$. In addition, it is known that if we choose all eigenvalues that are greater than some small ϵ, then there are about $2c/\pi$ eigenvalues from the flat region, $O(\log\left(\frac{1-\epsilon}{\epsilon}\right)\log c)$ eigenvalues from the transitional region, and $o(\log c)$ eigenvalues from the decay region (see [18] for a precise formulation). Thus, the number of eigenvalues greater then ϵ is dominated by the number of eigenvalues in the flat region, which is $2c/\pi$. As for the eigenvalues of the two-dimensional PSWFs, results in [29] indicate that as in the one-dimensional setting, the eigenvalues can be similarly divided into three distinct regions - flat, transitional, and super-exponential decay regions. Correspondingly, the number of significant eigenvalues is dominated by the number of eigenvalues in the flat region. Since $\lambda_{0,k}^{2c} \propto \alpha_{0,k}^{2c}$ (see (13)), it is clear that in order to satisfy condition (29) we need to determine the number of terms in the flat region of $\left|\lambda_{0,k}^{2c}\right|$. To this end, we follow the methods in [15] and [18] (see also [35] for the specific technique that we use), and consider the sum

$$\sum_{k=1}^{\infty}\left|\lambda_{0,k}^{2c}\right|^2, \tag{88}$$

which is approximately equal to the number of terms of $\left|\lambda_{0,k}^{2c}\right|$ which are close to 1 (denoted as the flat region). For simplicity of the presentation, we evaluate this sum for a bandlimit of c, and eventually, replace c with $2c$. By substituting $\varphi(r) = R(r)\sqrt{r}$ and $\gamma = \beta\sqrt{c}$ into (4), we obtain the integral equation

$$\gamma\varphi(r) = \int_0^1 \varphi(\rho)J_N(cr\rho)\sqrt{cr\rho}d\rho, \quad r\in[0,1], \tag{89}$$

whose eigenvalues and eigenvectors are $\varphi_{N,n}^c(r)$ and $\gamma_{N,n}^c$, respectively. Equation (89) was analysed in [31], where it is established that the eigenfunctions $\varphi_{N,n}^c(r)$ constitute a complete orthonormal system in $\mathcal{L}^2[0,1]$. Therefore, it follows that we have the identity

$$J_N(cr\rho)\sqrt{cr\rho} = \sum_{k=1}^{\infty}\gamma_{N,n}^c\varphi_{N,n}^c(r)\varphi_{N,n}^c(\rho), \tag{90}$$

for r and ρ both in $[0,1]$. If we notice that $\lambda^c_{N,n} = \sqrt{c}\gamma^c_{N,n}$, and take $N = 0$, we have

$$J_0(cr\rho)c\sqrt{r\rho} = \sum_{k=1}^{\infty} \lambda^c_{0,k}\varphi^c_{0,k}(r)\varphi^c_{0,k}(\rho). \tag{91}$$

Next, we take the squared absolute value of both sides of the equation above, followed by double integration (in r and ρ) to obtain

$$\int_0^1 \int_0^1 J_0^2(cr\rho)c^2 r\rho \, dr d\rho = \sum_{k=1}^{\infty} \left|\lambda^c_{0,k}\right|^2. \tag{92}$$

By evaluating the left hand-side of (92) using known integral identities of the Bessel functions [2], and after some manipulation, one can verify that

$$\sum_{k=1}^{\infty} \left|\lambda^c_{0,k}\right|^2 = \frac{c^2}{4}\left(J_0^2(c) - J_2(c)J_0(c) + 2J_1^2(c)\right). \tag{93}$$

Now, using the asymptotic approximation (see [2])

$$J_m(x) \sim \sqrt{\frac{2}{\pi x}}\cos(x - \frac{m\pi}{2} - \frac{\pi}{4}), \tag{94}$$

valid for $x \gg |m^2 - 1/4|$, we can write

$$\begin{aligned}\sum_{k=1}^{\infty} \left|\lambda^c_{0,k}\right|^2 &\sim \frac{c}{2\pi}\left(\cos^2(c - \frac{\pi}{4}) - \cos(c - \pi - \frac{\pi}{4})\cos(c - \frac{\pi}{4}) + 2\cos^2(c - \frac{\pi}{2} - \frac{\pi}{4})\right) \\ &= \frac{c}{2\pi}\left(2\cos^2(c - \frac{\pi}{4}) + 2\sin^2(c - \frac{\pi}{4})\right) \\ &= \frac{c}{\pi}. \end{aligned} \tag{95}$$

Therefore, the expression in (95) establishes that for large values of c, the number of terms in the set $\left\{\left|\lambda^{2c}_{0,n}\right|\right\}$ which are close to 1 is about $2c/\pi$ (see also Figure 4).

References

[1] Algorithms for single particle reconstruction. `http://spr.math.princeton.edu/`.

[2] Milton Abramowitz and Irene A Stegun. *Handbook of mathematical functions: with formulas, graphs, and mathematical tables*. Courier Corporation, 1964.

[3] Alok Dutt and Vladimir Rokhlin. Fast Fourier transforms for nonequispaced data. *SIAM Journal on Scientific computing*, 14(6):1368–1393, 1993.

[4] Mario Ferraro and Terry M Caelli. Relationship between integral transform invariances and Lie group theory. *JOSA A*, 5(5):738–742, 1988.

[5] Jeffrey A Fessler and Bradley P Sutton. Nonuniform fast Fourier transforms using min-max interpolation. *Signal Processing, IEEE Transactions on*, 51(2):560–574, 2003.

[6] Joachim Frank. *Three-Dimensional Electron Microscopy of Macromolecular Assemblies: Visualization of Biological Molecules in Their Native State*. Oxford, 2006.

[7] William T. Freeman and Edward H Adelson. The design and use of steerable filters. *IEEE Transactions on Pattern Analysis & Machine Intelligence*, 13(9):891–906, 1991.

[8] Leslie Greengard and June-Yub Lee. Accelerating the nonuniform fast Fourier transform. *SIAM review*, 46(3):443–454, 2004.

[9] Ran Hilai and Jacob Rubinstein. Recognition of rotated images by invariant Karhunen–Loéve expansion. *JOSA A*, 11(5):1610–1618, 1994.

[10] Matjaz Jogan, Emil Zagar, and Ales Leonardis. Karhunen–Loeve expansion of a set of rotated templates. *IEEE transactions on image processing: a publication of the IEEE Signal Processing Society*, 12(7):817–825, 2002.

[11] Kenichi Kanatani. *Group-theoretical methods in image understanding*, volume 20. Springer Science & Business Media, 2012.

[12] Jens Keiner, Stefan Kunis, and Daniel Potts. Using NFFT 3 – a software library for various nonequispaced fast Fourier transforms. *ACM Transactions on Mathematical Software (TOMS)*, 36(4):19, 2009.

[13] Boris Landa and Yoel Shkolnisky. Approximation scheme for essentially bandlimited and space-concentrated functions on a disk. *Applied and Computational Harmonic Analysis*, 2016.

[14] Henry J. Landau. Necessary density conditions for sampling and interpolation of certain entire functions. *Acta Mathematica*, 117(1):37–52, 1967.

[15] Henry J. Landau. On szegös eingenvalue distribution theorem and non-hermitian kernels. *Journal d'Analyse Mathématique*, 28(1):335–357, 1975.

[16] Henry J Landau and Henry O Pollak. Prolate spheroidal wave functions, Fourier analysis and uncertainty - II. *Bell System Technical Journal*, 40(1):65–84, 1961.

[17] Henry J Landau and Henry O Pollak. Prolate spheroidal wave functions, Fourier analysis and uncertainty - III: The dimension of the space of essentially time-and band-limited signals. *Bell System Technical Journal*, 41(4):1295–1336, 1962.

[18] Henry J. Landau and Harold Widom. Eigenvalue distribution of time and frequency limiting. *Journal of Mathematical Analysis and Applications*, 77(2):469–481, 1980.

[19] Catherine L. Lawson, Ardan Patwardhan, Matthew L. Baker, Corey Hryc, Eduardo Sanz Garcia, Brian P. Hudson, Ingvar Lagerstedt, Steven J. Ludtke, Grigore Pintilie, Raul Sala, John D. Westbrook, Helen M. Berman, Gerard J. Kleywegt, and Wah Chiu. Emdatabank unified data resource for 3dem. *Nucleic Acids Research*, 44(D1):D396–D403, 2016.

[20] Reiner Lenz. Group-theoretical model of feature extraction. *JOSA A*, 6(6):827–834, 1989.

[21] Reiner Lenz. Group invariant pattern recognition. *Pattern Recognition*, 23(1):199–217, 1990.

[22] Reiner Lenz. *Group theoretical methods in image processing*. Springer-Verlag New York, Inc., 1990.

[23] Satya P. Mallick, Bridget Carragher, Clinton S. Potter, and David J. Kriegman. ACE: automated CTF estimation. *Ultramicroscopy*, 104(1):8–29, 2005.

[24] Andrei Osipov, Vladimir Rokhlin, and Hong Xiao. Prolate spheroidal wave functions of order zero. *Springer Ser. Appl. Math. Sci*, 187, 2013.

[25] Pietro Perona. Deformable kernels for early vision. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 17(5):488–499, 1995.

[26] Daniel P Petersen and David Middleton. Sampling and reconstruction of wave-number-limited functions in n-dimensional euclidean spaces. *Information and control*, 5(4):279–323, 1962.

[27] Colin Ponce and Amit Singer. Computing steerable principal components of a large set of images and their rotations. *Image Processing, IEEE Transactions on*, 20(11):3051–3062, 2011.

[28] Daniel Potts, Gabriele Steidl, and Manfred Tasche. Fast Fourier transforms for nonequispaced data: A tutorial. In *Modern sampling theory*, pages 247–270. Springer, 2001.

[29] Kirill Serkh. On generalized prolate spheroidal functions. Technical Report TR-1519, Department of Mathematics, Yale University, 2015.

[30] Yoel Shkolnisky. Prolate spheroidal wave functions on a disc - integration and approximation of two-dimensional bandlimited functions. *Applied and Computational Harmonic Analysis*, 22(2):235–256, 2007.

[31] David Slepian. Prolate spheroidal wave functions, Fourier analysis and uncertainty - IV: extensions to many dimensions; generalized prolate spheroidal functions. *Bell System Technical Journal*, 43(6):3009–3057, 1964.

[32] David Slepian and Henry O Pollak. Prolate spheroidal wave functions, Fourier analysis and uncertainty - I. *Bell System Technical Journal*, 40(1):43–63, 1961.

[33] Mark Tygert. Recurrence relations and fast algorithms. *Applied and Computational Harmonic Analysis*, 28(1):121–128, 2010.

[34] Cédric Vonesch, Frédéric Stauber, and Michael Unser. Steerable PCA for rotation-invariant image recognition. *SIAM Journal on Imaging Sciences*, 8(3):1857–1873, 2015.

[35] Hong Xiao, Vladimir Rokhlin, and Norman Yarvin. Prolate spheroidal wavefunctions, quadrature and interpolation. *Inverse problems*, 17(4):805, 2001.

[36] Zhizhen Zhao, Yoel Shkolnisky, and Amit Singer. Fast steerable principal component analysis. *IEEE Transactions on Computational Imaging*, 2(1):1–12, 2016.

[37] Zhizhen Zhao and Amit Singer. Fourier–Bessel rotational invariant eigenimages. *JOSA A*, 30(5):871–877, 2013.